

Communication Requirements and Informative Signaling in Matching Markets

Itai Ashlagi Mark Braverman Yash Kanoria Peng Shi*

July 2, 2017

Abstract

We study how much communication is needed to find a stable matching in a two-sided matching market with private preferences. Segal (2007) and Gonczarowski et al. (2015) showed that, in the worst case, any protocol that computes a stable matching requires the communication cost per agent to scale linearly with the total number of agents. In markets with many thousands of agents, this communication requirement is implausibly high, casting doubt on whether stable matchings can arise in large markets.

We study markets with realistic assumptions on the preferences of agents and their available information, and show that a stable matching can be found with a much smaller communication requirement. In our model, the preferences of workers are unrestricted, and the preferences of firms follow an additively separable latent utility model. Our efficient communication protocol modifies the worker-proposing deferred acceptance algorithm by having firms signal workers they especially like while also broadcasting qualification requirements to discourage other workers who have no realistic chances from applying. In the special case of tiered random markets, the protocol

*Ashlagi: Department of Management Science and Engineering, Stanford University, iashlagi@stanford.edu. Braverman: Computer Science, Princeton University, mbraverm@cs.princeton.edu. Kanoria: Graduate Business School, Columbia University, kanoria@columbia.edu. Shi: Department of Data Sciences and Operations, Marshall School of Business, University of Southern California, pengshi@usc.edu. This paper has greatly benefited from conversations with Avinatan Hassidim, Ilya Segal, and Assaf Romm. We also thank Federico Echenique, Marina Halac, Scott Kominers, Anqi Li, Jacob Leshno, and Marek Pycia for very helpful comments. Ashlagi acknowledges the research support of the National Science Foundation (grant SES-1254768).

can be modified to run in two rounds and involve only private messages. Our protocols have good incentive properties and give insights into how to mediate large matching markets to reduce congestion.

1 Introduction

Stable matching is a widely used solution concept for two-sided matching markets. In a stable matching, no pair of agents would prefer to match with each other over their assigned partners.¹ This solution concept has been important in shaping the design of real matching markets such as the National Resident Matching Program (NRMP),² as well as school choice programs in many cities in the United States.³ Stability has also been the identifying assumption in many empirical studies of structural estimation in matching markets.⁴

A natural question is, How do stable matchings form? The current answer in the literature is as follows. In centralized markets, the deferred acceptance (DA) algorithm can be used to find a stable matching (Gale and Shapley, 1962). In this algorithm, one side of the market *applies* to the other side, which tentatively accepts the applicant it prefers most and rejects all others. The rejected applicants then apply to their next choices, and the algorithm iterates until convergence. In decentralized markets, a sequence of proposals and responses can converge to a stable matching under certain assumptions, but the time required may be large.⁵

However, the above answer has a deficiency in its logic: it does not account for constraints on communication and information acquisition. In centralized markets, DA is guaranteed to yield a stable outcome only if agents submit complete rankings of eligible partners. However, real systems have constraints on how long an agent’s preference list can be. For example,

¹See Roth and Sotomayor (1990) for an overview of the theory of stable matching.

²See Roth (1984), Roth and Peranson (1999), and Roth (2002).

³See Abdulkadiroğlu et al. (2005a), Abdulkadiroğlu et al. (2005b), and Abdulkadiroğlu et al. (2009).

⁴Empirical works include both those with transferable utility (Choo and Siow, 2006; Galichon and Salanié, 2015), as well as those with nontransferable utility (Hitsch et al., 2010; Sørensen, 2007; Agarwal, 2015; Galichon et al., 2016). See Chiappori and Salanié (2016) for a review of both literatures. This paper focuses on the nontransferable case, although the two are mathematically related (Kelso and Crawford, 1982; Hatfield and Milgrom, 2005).

⁵See Roth and Vate (1990), Adachi (2003), and Ackermann et al. (2011) for the convergence to stable matching in decentralized markets. Other papers that study the conditions for such convergence include Pais (2008), Niederle and Yariv (2009), Haeringer and Wooders (2011), and Echenique and Yariv (2013).

the New York City NYC school choice system allows students to report preference rankings of length at most 12. In the NRMP, medical students cannot rank a hospital among their preferences unless they interview there, and the logistical cost of interviews limits the lengths of rankings. In decentralized markets, the outcome may not be a stable matching precisely because the market cannot process communications fast enough. This market congestion has been empirically studied in laboratory experiments and in the field.⁶

Recent theoretical results suggest that it is indeed difficult for large markets to arrive at stable matchings. Using the theory of communication complexity,⁷ which studies the minimum communication required to accomplish certain tasks from an information theory perspective, Segal (2007), Chou and Lu (2010), and Gonczarowski et al. (2015) prove that for any method of finding stable matchings, there exists a distribution of preferences in which agents must *learn* and *communicate* their preferences for a substantial fraction of the entire market. More precisely, the worst-case amount of communication per agent must grow linearly with the number of agents.⁸ This communication requirement is implausibly large for many real markets, which have many thousands of agents on each side.⁹ This communication overhead makes the solution concept of stability less plausible in large markets.

The main contribution of this paper is to begin correcting this deficiency in the theory by showing that under natural assumptions about the distribution of preferences and prior knowledge of agents, stable matchings *can form* under limited communication in large markets, if agents participate in informative signaling. By “informative,” we mean that certain signals are more useful than others in helping the market converge to a stable match. The idea is to construct signals that help agents to estimate whom they can realistically be

⁶See Roth and Xing (1994) and Avery et al. (2001) for empirical findings of unravelling in labor markets, Kagel and Roth (2000) for experiments towards reducing unravelling, and Roth (2008) for a brief survey on congestion in matching markets.

⁷See Kushilevitz and Nisan (2006) for a review of the communication complexity literature. The importance of studying communication in economic models is highlighted in the seminal essay Hayek (1945). This research direction was first formalized in Hurwicz (1973) and Mount and Reiter (1974).

⁸Segal (2007) proves that the communication per agent may be linear in the number of agents, assuming deterministic preferences, deterministic communication protocol, and exact stability. Chou and Lu (2010) extend this to approximate stability. Gonczarowski et al. (2015) extend it to randomized preferences and protocols. We give a precise restatement in Section 2.3.1.

⁹In 2017, the number of applicants in the NRMP is about 43,000, and the number of positions is 31,000 (National Residency Matching Program, 2017). The American labor market has about 160 million workers (US Bureau of Labor Statistics, 2017).

matched with and to encourage agents to reach out to easy-to-get partners, while waiting for harder-to-get partners to reach out to them. We show that when every agent participates in such signaling, the market can with high probability reach a stable matching with low levels of communication and preference learning. Moreover, under certain assumptions, it is in the best interest of each agent to comply with signaling scheme, assuming that others also do so.

Before describing our results in detail, we first comment on how this paper is conceptually distinct from previous works on signaling and information friction in matching markets.¹⁰ First, the role of signaling in this paper differs, in that we focus on the *informative* aspect of signaling rather than the *strategic* aspect. In practical terms, in previous work, signaling changes the set of equilibrium outcomes. In this paper, signaling does not change the outcome but enables the market to reach this outcome more quickly. The earlier perspective stems from the literature on signaling games in economics, whereas the perspective in this paper stems from information theory. Both interpretations have relevance for real markets. This paper bridges these two literatures and shows that signals are useful not only to improve the matching outcome but also to speed up communication.

Second, our approach on modeling information friction differs from previous approaches based on search theory, in which agents optimize the matching outcome given costs or constraints on communication or information acquisition.¹¹ On the contrary, we take the dual approach, in which the matching outcome is fixed, and we look for the least amount of communication needed to obtain that outcome. These approaches complement each other. The former focused on the effect of search cost on the matching outcome; our approach focuses on how to make the search more efficient, given a desired outcome.¹²

We now describe our precise assumptions and results. Our assumptions on the preference distribution are mild. We say that a market is *separable* if the preferences of one side, say, firms, follow a latent random utility model with an additively separable structure. The preferences of workers can be arbitrary. For a firm, its latent utility for a worker is the sum

¹⁰See Coles et al. (2010), Chade et al. (2014), Lee and Niederle (2015), Abdulkadiroğlu et al. (2015), and Lee and Schwarz (2017).

¹¹See Eeckhout (1999), Adachi (2003), Das and Kamenica (2005), Lauer mann and Nöldeke (2014), Arnosti et al. (2014), Kanoria and Saban (2017). For a review of the use of search theory in matching, see Chade et al. (2014).

¹²Other studies that focus on making search more efficient include Arnosti et al. (2014), Kanoria and Saban (2017), and Horton (2017). However the outcome of interest is not a stable matching in these works.

of a public score and a private score, both of which are heterogeneous across firm-worker pairs. The public score represents the worker’s general level of fit based on observable characteristics such as past experience, level of education, and test scores. This information is known to the worker and the firm, but not to anyone else. The private score represents the idiosyncratic component of the firm’s preference and is drawn independently from a certain unknown distribution. The preferences of agents, including both workers and firms, are unknown a priori to everyone, even to the agents themselves. In order to learn their own preferences, agents have to query a choice function, which returns their most preferred partner within a given set.

Because the public scores can be heterogeneous across worker-firm pairs, we need the private scores to be sufficiently important relative to the public scores; otherwise, separable markets would include general markets, and the previous impossibility results of Gonczarowski et al. (2015) would hold. Precisely speaking, we assume that the range of public scores and the hazard rate of private scores are bounded above. A bounded hazard rate allows the private scores to take any heavy-tailed distribution, including the exponential, type-I extreme value, lognormal, and Pareto distributions.

Our main result (Theorem 1) is that in any separable market, there exists a way to find a stable matching with high probability using low levels of communication and preference learning. The protocol that we construct, called communication-efficient deferred acceptance (CEDA), modifies the worker-proposing DA algorithm by having workers apply only to firms where they have a realistic chance. Workers know whether they have a chance through signals sent by firms. We show that in any separable market, this protocol yields, with high probability, the worker-optimal stable match. Furthermore, the communication cost, as measured by the number of bits agents send on average, and the preference learning cost, as measured by the number of choice function queries agents make on average, both scale according to the square root of the market size,¹³ which is much lower than the linear scaling necessary under arbitrary preferences.

CEDA’s efficiency lies in the signals sent by firms, which help workers direct their applications. There are two types of signals: in the beginning, each firm partitions workers into different sets based on their public scores and sends a *preference signal* to a certain number of its favorite workers in each set. Receiving a preference signal indicates to a worker that

¹³The square root comes from the communication complexity of set disjointness, which is studied in Babai et al. (1986), Kalyanasundaram and Schnitger (1992), Razborov (1992), and Braverman et al. (2013).

she should apply despite having a potentially unfavorable public score. During the second, worker-proposing phase of CEDA, each firm maintains and updates a public *qualification requirement*, which is broadcast to all workers. The qualification requirement is increased whenever the firm receives sufficiently many applications from workers whose public score meets the requirement.¹⁴ We show that if a worker neither meets the qualifying public score nor receives a preference signal from a firm, then her chance of being matched with that firm is essentially zero, so she should not waste her time applying. Moreover, when workers apply only to firms where their public scores qualify or from which they have received preferences signals, then the total communication cost is low, with high probability.

For markets with tens of thousands of agents, CEDA requires each agent to have hundreds of interactions,¹⁵ which, one may argue, is a reasonable number. A natural question is whether a more efficient protocol exists. Our second result (Theorem 2) is that the answer is essentially no. We give a simple example of a separable market in which any protocol that finds a stable matching with high probability must incur communication and preference learning costs that scale according to the square root of the market size. Therefore, the CEDA protocol is near-optimal for separable markets in terms of its efficiency of communication.

However, the CEDA protocol may require many rounds of communication, as what signal an agent should send may depend on what signals the agent has received, which may in turn depend on the prior signals of other agents. We give examples that illustrate that such sequential dependency may be necessary for any protocol under certain preference distributions. However, given stronger assumptions on preferences, we show that a two-round protocol is possible. The markets we consider are *tiered random markets*, in which agents on both sides are partitioned into tiers; an agent of a higher tier is always preferable to one of a lower tier, and preferences are uniformly random and independent within each tier.¹⁶ For such markets, we give a two-round protocol whose preference learning and communication costs scale only polylogarithmically in the market size (Theorem 3). The protocol has the additional advantage of using only private signals, which need to be seen only by a single receiver. The protocol, called the *targeted-signaling protocol*, designates for each agent a set of

¹⁴We show that when the qualification requirement increases, it would turn away only applicants with negligible chances of being accepted anyway. The reason is that preferences in separable markets are sufficiently idiosyncratic, so after receiving a certain number of applications from qualified applicants, the firm would most likely have found a worker it sufficiently likes.

¹⁵The reason is that the communication requirement of CEDA scales as the square root of the market size.

¹⁶Such preferences are called “block correlated” in Coles et al. (2013).

easy-to-get partners and hard-to-get partners, based on commonly known tier information alone. In the first round, each agent signals a certain number of her favorite easy-to-get partners. In the second round, each agent submits a partial preference ranking, in which she ranks only the subset of potential partners whom she signaled or who signaled her. The protocol outputs a matching based on these partial preferences, and we show that this matching is stable with respect to the full preferences with high probability.

Both our protocols inherit the incentive properties of worker-proposing DA under complete information. More precisely, we show that with high probability, no worker can unilaterally deviate from the protocol and improve her outcome, and no firm can unilaterally deviate and be matched to someone better than its partner in the firm-optimal stable match (Theorem 4). Recent literature has demonstrated that in large markets, under mild assumptions, the vast majority of agents may have the same match partner under the worker-optimal and firm-optimal stable matchings.¹⁷ For such markets, all agents have vanishing incentives to deviate from our protocol.

Our results suggest that stable matching is a plausible solution concept for real markets, even when the market is large and communication is limited. A supporting observation is that many markets do use the types of signals in our protocols, and our theory highlights the importance of such signals for achieving efficient communication. For example, in the labor market, firms post requirements for years of experience, education, and other observable qualities, and these can play the role of qualification requirements in CEDA in guiding the applications of workers. Similarly, certain universities in Israel, India, and Iran publish predicted admission thresholds, which can help students better target their applications. Moreover, firms also reach out to potential applicants to indicate interest, either directly or through headhunters, which is similar to preference signaling in our model. Our protocols also help to explain observed behaviors in matching markets. For example, in laboratory experiments of deferred acceptance, proposing agents skip potential partners for whom they have low chances of matching (Echenique et al., 2016). Moreover, empirical evidence on signaling in the economics job market suggests that candidates signal departments that are in tiers lower than or similar to their own (Coles et al., 2010), mirroring the structure of the targeted-signaling protocol.¹⁸

¹⁷See Immorlica and Mahdian (2005), Kojima and Pathak (2009), Ashlagi et al. (2017), and Lee (2017).

¹⁸Coles et al. (2010) only document the formal signaling provided by the American Economics Association, but our notion of preference signaling also corresponds to the informal signaling as well.

Besides illustrating how stable matchings can form in large markets under limited communication, our results also provide prescriptive guidelines for market design.¹⁹ A common feature of both protocols is that they estimate for each agent the set of easy-to-get partners and hard-to-get partners, and recommend agents to reach out to their favorite easy-to-get partners while waiting for hard-to-get partners to reach out to them. In centralized clearing-houses such as the NRMP, the market designer can adapt this idea to give agents guidelines for targeting their search in the pre-market. In online matching platforms for dating and labor, the platform can use past data to estimate levels of fit and who the hard-to-get partners might be, and they can apply insights from our protocols to improve the overall efficiency of search. We discuss these recommendations more concretely in Section 6.

2 Model

2.1 Review of asymptotic notation

This paper heavily utilizes Big-O type asymptotic notation. For clarity, we define these notations in this section.

Given two nonnegative functions $f, g : \mathbb{N} \rightarrow \mathbb{R}_+$, we say that $f(n) = O(g(n))$ if there exists n_0 and $M > 0$ such that $f(n) \leq Mg(n)$ for all $n \geq n_0$. This indicates that up to a multiplicative constant, the tail of the function f grows no faster than g .

We say that $f(n) = o(g(n))$ if for any $\epsilon > 0$, there exists n_0 such that $f(n) \leq \epsilon g(n)$ for all $n \geq n_0$. A special case is $f(n) = o(1)$, which means that $\lim_{n \rightarrow \infty} f(n) = 0$.

We say that $f(n) = \Omega(g(n))$ if $g(n) = O(f(n))$. In other words, there exists n_0 and $M > 0$ such that $f(n) \geq Mg(n)$ for all $n \geq n_0$.

Finally we say that $f(n) = O^*(g(n))$ if there exists n_0 and $C > 0$ such that $f(n) \leq (\log n)^C g(n)$ for all $n \geq n_0$.

2.2 Two-sided matching markets with preference learning

In this section, we define a generic model of two-sided matching markets with incomplete information. The markets studied in this paper, separable markets (Section 2.4) and tiered random markets (Section 4.1), are special cases of this more general model. Compared to

¹⁹For other recent studies on improving the efficiency of search in online markets, see Horton (2017), Fradkin (2017), and Kanoria and Saban (2017).

previous work, the model has two distinguishing features: agents are allowed to have partial information on the preference distribution of others, and agents do not know their own preferences directly, but must query a choice function to learn them.

A *two-sided matching market* \mathcal{M} is defined by a tuple $(I, J, \omega, K, \mathcal{P})$, where $I = \{1, 2, \dots, n_I\}$ is a set of workers, and $J = \{n_I + 1, \dots, n_I + n_J\}$ is a set of firms. Both workers and firms are called *agents*. For concreteness, we refer to each worker as “she” and each firm as “it.” Let $n = n_I + n_J$. This is the total number of agents, and we call this the *market size*. The parameter $\omega \in W$ represents the true *state of the world*, where W is the set of possible states. K represents the *a priori knowledge* of agents about the state of the world. It is indexed by each agent $a \in I \cup J$, and K_a is a subset of W that is guaranteed to contain the true state of the world ω . A *preference realization* R is indexed by the set of agents $I \cup J$. For worker $i \in I$, R_i is a permutation of $J \cup \{0\}$, which specifies her (strict) *preference ordering* over being matched with each firm and remaining unmatched, which is represented by the symbol 0. Similarly, for each firm $j \in J$, R_j is a permutation of $I \cup \{0\}$, which specifies the *preference ordering* of firm j . A potential partner is said to be *acceptable* to an agent if the agent prefers being matched to the partner over being unmatched. The function $\mathcal{P} : W \rightarrow \Delta(R)$ is called the *preference function*, and maps the state of the world to a probability distribution over preference realizations. Since ω is the true state of the world, $\mathcal{P}(\omega)$ is called the *true preference distribution*. We assume that each agent $a \in I \cup J$ knows a priori that $K_a \ni \omega$ and the function \mathcal{P} , but not the true state of the world ω and not the realization R_a for the agent’s own preference ordering.

To learn the realization of their own preference ordering, each agent must query their *choice function*, which specifies the agent’s most preferred option within a given set of potential partners. Precisely speaking, for each worker $i \in I$, the choice function is $C_i : 2^{J \cup \{0\}} \rightarrow J \cup \{0\}$, which takes as input a subset of possible options, $S \subseteq J \cup \{0\}$, and outputs the highest-ranked element according to the preference ordering R_i , which the agent does not directly observe. Similarly, the choice function is defined for every firm $j \in J$, except the set of possible options is now $I \cup \{0\}$. Later we quantify the preference learning cost of an agent in terms of the number of choice function queries the agent must execute during the market clearing process.

2.2.1 Stable matchings

Having defined two-sided matching markets, we define stable matchings. A collection of worker-firm pairs $\mu \subseteq I \times J$ is called a *matching* if each worker and each firm appears at most once. We refer to each $(i, j) \in \mu$ as a *matched pair*, and i and j as *matched partners* to one another in μ . Agents who have no matched partner are said to be *unmatched* in μ . While a matching μ is technically speaking a set of tuples, we abuse notation and also use μ as a function that maps each agent to their matched partner: if (i, j) is a matched pair, then $\mu(i) = j$ and $\mu(j) = i$; if an agent $i \in I \cup J$ is unmatched, then $\mu(i) = 0$. A potential partner is said to be *acceptable* to an agent if the agent prefers to be matched to the partner over being unmatched. A *blocking pair* to a matching μ is a pair $(i, j) \in I \times J$ which is not a matched pair, but both agents find one another acceptable, and worker i prefers firm j to $\mu(i)$, and firm j prefers worker i to $\mu(j)$. A matching is said to be *stable* if there are no blocking pairs. For any preference realization R , a stable matching always exists and can be found using the deferred acceptance (DA) algorithm of Gale and Shapley (1962). A version of this is formally stated in Section 2.3.

2.3 Stable matching protocol and communication cost

In this section, we formalize the concept of the communication cost of finding a stable matching, using concepts from the *communication complexity* literature. (See Kushilevitz and Nisan, 2006, for an overview of this literature.)

We first define a *communication protocol* as follows. There is a set of agents, each of whom can send messages to other agents. Each message is formally represented as a sequence of zero-one bits, and we measure the length of each message by the number of bits. For now, we assume that messages are *public*, which means they are visible to all agents. (We remove this assumption in Section 4, where we study protocols in which messages are visible only to a particular receiver.) We define the *history* of messages to be the sequence of all messages sent by any agent since the beginning of the protocol. As a possibly randomized function of the current history of messages, the protocol either terminates with a final output or chooses one agent to send the next message. The protocol also specifies what the next agent's message should be as a possibly randomized function of the agent's private information and the history of messages.

We emphasize the following definitions, which are needed to state our main results.

Definition 2.1. Given a two-sided matching market $\mathcal{M} = (I, J, \omega, K, \mathcal{P})$, a matching protocol Π is a communication protocol in which the set of agents is $I \cup J$ and the output is a matching $\mu \subseteq I \times J$. A **matching protocol** Π is said to be **stable with high probability** in market \mathcal{M} if the probability that the matching produced is stable with respect to the preference realization R converges to 1 as the market size n goes to ∞ . Here, the probability is taken over the distribution $\mathcal{P}(\omega)$ of preferences as well as any randomness in the protocol.

Definition 2.2. Given a two-sided matching market $\mathcal{M} = (I, J, \omega, K, \mathcal{P})$ and a matching protocol Π , the **communication cost** is the expected total length of all messages (in bits) divided by the total number of agents. The **preference learning cost** is the expected total number of choice function queries, divided by the total number of agents. The expectations in both definitions above are defined over the distribution $\mathcal{P}(\omega)$ of preferences as well as any randomness in the protocol.

An example of a matching protocol that is always stable is the DA algorithm, given below. (In particular, this is the sequential version of the worker-proposing DA algorithm, developed by McVitie and Wilson (1971).)

Protocol 1. *The sequential deferred acceptance (DA) algorithm.*

The protocol keeps track of a tentative matching, which is initialized to be empty.

1. *Consider the set of tentatively unmatched workers who have not yet applied to all firms they find acceptable. If this set is empty, then the algorithm terminates and outputs the current tentative matching. Otherwise, one worker from this set is selected arbitrarily, say worker i , and she applies to her favorite firm j that she finds acceptable to which she has not yet applied.*
2. *If firm j is tentatively unmatched and finds her acceptable, then it becomes tentatively matched to the worker. Otherwise, if it is already tentatively matched to some other worker, say worker i' , then it becomes tentatively matched to the more preferred worker among the two and rejects the other. The rejected worker becomes tentatively unmatched and can again be chosen in a future step. Return to step 1.*

This protocol requires high preference learning and communication cost: each step requires a choice function query from an agent, as well as a message of length $\Omega(\log n)$ to encode an application or a response. Hence, in the worst case, when each worker applies to

every firm, the communication cost is $\Omega(n \log n)$ per agent and the preference learning cost is $\Omega(n)$ per agent.

2.3.1 Impossibility of sublinear communication cost for arbitrary markets

The following negative result, which is a strengthening of an earlier result²⁰ by Segal (2007), implies that without restrictions on the preference distribution, one cannot hope to improve much upon the DA algorithm. There exists a distribution of preferences such that any matching protocol that is stable with high probability must require agents to learn and communicate on average at least a constant fraction of their preferences over the entire market. For real world markets with many thousands of agents on each side, this is an unrealistically high requirement for communication and preference learning.

Proposition 2.3. *(adapted from Gonczarowski et al., 2015) There exists a two-sided matching market $\mathcal{M} = (I, J, \omega, K, \mathcal{P})$ with $K_i = \{\omega\}$ (everyone knows the state of the world and hence the precise distribution of preferences), such that any matching protocol that is stable with high probability requires a communication cost of at least $\Omega(n)$ per agent and a preference learning cost of at least $\Omega(\frac{n}{\log n})$ per agent.²¹*

However, this negative result requires a fairly contrived distribution of preferences. The takeaway from this paper is that since preferences in real markets are not worst-case, but usually exhibit additional structure, this structure can be used to find stable matchings much more efficiently than what the above result implies.

2.4 Separable markets

In this section, we define separable markets, a restricted class of two-sided matching markets for which we show in Section 3 that the communication requirement is much less than in

²⁰The earlier result of Segal (2007) implies that for any matching protocol that always produces a stable matching, there exists a preference realization R such that the protocol requires an average communication cost of $\Omega(n)$ per agent.

²¹The original results in Gonczarowski et al. (2015) concern communication cost only and preference learning cost. However, any lower bound on communication cost automatically implies a bound on preference learning, because any protocol that uses Q choice function queries can be made into a communication protocol with $Q \log n$ bits of communication, because the only information relevant to computing a stable matching is the result of choice function queries.

arbitrary markets. After precisely stating the assumptions behind separable markets, we discuss in Section 2.4.1 why this is a reasonable model of matching markets.

A separable market $\mathcal{M} = (I, J, \omega, K, \mathcal{P})$ is a two-sided matching market with the following restrictions on the preference function \mathcal{P} , the true state of the world ω , and the a priori knowledge K of the agents.

We assume no restrictions on the preferences of one side. Without loss of generality, let the unrestricted side be the workers. We assume the preferences of the firms follow a latent random utility model with the following additively separable structure. (Hence the term “separable” markets.) The latent utility of firm j for worker i is

$$u_{ji} = a_{ji} + \epsilon_{ji}, \quad (1)$$

where a_{ji} is the *public score* of worker i for firm j , which represents the *observable* characteristics for this worker-firm pair, and ϵ_{ji} is the *private score* of worker i for firm j , which represents the *idiosyncratic* component of the firm’s preference for that worker. The utility of the firm for being unmatched, u_{j0} , is unrestricted. Denote this as firm j ’s *outside utility*. We assume that all private scores for firm j are distributed independently and identically according to a distribution F_j and are independent of public scores, the preferences of workers, and the utility of firms for being unmatched.

Let $\mathcal{P}(\omega)$ denote the preference distribution described above, where the state of the world ω encapsulates the preferences of workers and the outside utilities of firms, both of which are completely unrestricted, as well as the set of public scores $\{a_{ji}\}$ and private score distributions $\{F_j\}$. These latter two satisfy the following parametric assumptions.

Assumption 2.4 (range of public scores). *For every firm j , its range of public scores is upper-bounded by a polylogarithmic function of market size n :*

$$\bar{a}_j - \underline{a}_j = O^*(1), \quad (2)$$

where $\bar{a}_j = \max_{i \in I} \{a_{ji}\}$, and $\underline{a}_j = \min_{i \in I} \{a_{ji}\}$.²²

Assumption 2.5 (bounded hazard rate). *For every firm j , the distribution of private scores, F_j , satisfies a uniform bound on its hazard rate:*

$$h(x) := \frac{F'_j(x)}{1 - F_j(x)} \leq 1 \quad \forall x \in \mathbb{R}. \quad (3)$$

²²In other words, there exists a constant $C > 0$ such that $\bar{a}_j - \underline{a}_j \leq \log^C n$.

The bound in Assumption 2.5 is normalized to 1 without loss of generality, because one can multiplicatively scale latent utilities by an arbitrary positive constant without affecting the underlying preferences. A bounded hazard rate is exhibited by any distribution with a sufficiently heavy tail, including the exponential, the type-I extreme value, the log-normal, and the Pareto distributions.

Assumptions 2.4 and 2.5 together guarantee that the idiosyncratic component of firm preferences (private scores) is sufficiently important compared to the systematic component (public scores). As will be further explained in Section 2.4.1, these assumptions still allow a substantial systematic component, and some version of these assumptions is *necessary* to bypass the impossibility result of Proposition 2.3.

The assumptions on the a priori knowledge K of agents are as follows. Each agent knows the public scores that associated with them. In particular, the public score a_{ji} is a priori known to firm j and worker i , but not to anyone else. In addition, workers and firms know that the state of the world ω satisfies Assumptions 2.4 and 2.5, but they do not know the specific bound on the hazard rate or the specific distribution F_j for any firm. They also do not know anything about the other components of the state of the world, which are the preferences of workers and the outside utility of firms.

Beside the a priori knowledge K , each agent can learn their own preference realization by querying their own choice function, as described in Section 2.2. Agents can learn about the other's preferences only through communication.

2.4.1 Discussion of assumptions

In this section, we justify the key assumptions behind separable markets, which are:

1. Additive separable latent utilities for firms, with public score a_{ji} mutually known to firm j and worker i .
2. Independent private scores, drawn from the same distribution for each firm.
3. Bounds on the range of public scores and on the hazard rate of the private score distributions (Assumptions 2.4 and 2.5).

The additive separable structure of preferences for one side of the market is motivated by the following observation: in many real matching markets, the preference of at least one side depends on observable characteristics in a predictable way. For example, in hiring for a

position, a firm may value an applicant’s education, GPA, relevant certification, and relevant work experience. A worker potentially interested in applying may also have reasonable a priori knowledge of how important each of these characteristics is for a particular position, and can assess her general level of fit without communicating with the firm. Likewise, without the need for communication, the firm may observe many of a potential worker’s characteristics from LinkedIn or a university’s alumni database, and the worker may have a priori information about the firm or about the industry. The public score a_{ji} represents this mutually observable general level of fit between worker i and firm j . Note that we allow this to vary for each worker-firm pair, allowing for rich heterogeneities. The private score ϵ_{ji} represents everything that is unexplained by the observable component, including for example a firm’s unobserved preference for a particular type of background or skill that a worker may not foresee, or the inherent idiosyncrasies in firm’s perception of workers.

The assumption that private scores are independent and commonly distributed is for technical convenience. This assumption is prevalent in the discrete choice literature as well as in most empirical studies in matching markets.²³ In our analysis, this assumption allows us to claim that with high probability, after examining many workers, the firm must have found someone with a high private score. Furthermore, it allows us to bound the number of workers with a private score above a certain quantile. Because these are the only times we use this assumption, we expect the analysis to be generalizable to models in which the unobservable component of firm preferences exhibits mild correlations or mild variations in magnitude across workers for each firm.²⁴

Our assumptions on the existence of bounds for the range of public scores and on the hazard rate can be interpreted as a condition that the idiosyncratic component of firm preference is sufficiently important relative to the systematic component. These are perhaps the most restrictive of our assumptions. Regarding these assumptions, we make three observations.

First, some type of such assumptions are necessary to bypass the previous impossibility results on communication-efficient stable matching protocols (see Proposition 2.3). For example, if private scores were identically zero, or if the public scores were to have an unbounded range, then one can make the public scores the only relevant component in the firm

²³See for example Choo and Siow (2006), Agarwal (2015), Dagsvik (2000), Peski (2015), Menzel (2015), and Galichon et al. (2016), as well as the review article Chiappori and Salanié (2016).

²⁴However, if the correlations in preferences are allowed to be large and complex, then the negative result of Gonczarowski et al. (2015) applies and no communication efficient protocol exists.

preferences. In this case, one can embed any arbitrary deterministic preferences into the separable market model, because the public scores $\{a_{ji}\}$ are completely flexible. Furthermore, the fact that worker i knows a_{ji} for every firm does not help the worker, because the worker cannot meaningfully compare public scores across firms, as the additive scale for each firm may be arbitrary. In this case, we are back to the setting of private, arbitrary preferences on both sides, so previous impossibility results would apply.

Second, these assumptions still allow a substantial amount of systematic variation in the preferences of each firm. For example, suppose F_j is the exponential distribution with parameter 1 (which satisfies Assumption 2.5). In this case, a difference in public scores of $3 \log n$ implies that the firm will prefer the worker with the higher public score with a probability of $1 - \frac{1}{n^3}$, which remains high even after a union bound over the $O(n^2)$ worker-firm pairs. Thus, our assumptions allow enough systematic variation for each firm to have multiple tiers of workers, such that, with high probability, every worker in a better tier is preferred to every worker in a lower tier. The number of tiers can also grow to infinity, as long as it is controlled by a polynomial of $\log n$. Moreover, we note that these tiers may be the same for different firms or distinct for different firms.²⁵

Finally, these assumptions can be relaxed and our results can be generalized. In Appendix A.1, we show how the $O^*(1)$ bound on the range of public scores can be generalized to a bound of $g(n)$ for a general sublinear function $g(\cdot)$.

3 Main results

The main results for separable markets are as follows.

1. For separable markets, there exists a matching protocol that is stable with high probability and that in the worst case incurs communication and preference learning costs of $O^*(\sqrt{n})$ per agent (see Theorem 1). This is much lower than the $\Omega(n)$ cost needed for arbitrary two-sided markets (see Proposition 2.3). The protocol gives insights into what the informative signals are.
2. This $O^*(\sqrt{n})$ communication and preference learning costs is essentially the best possible guarantee for separable markets, as we give an example of a separable market in

²⁵In comparison, the preference distributions in Kojima and Pathak (2009) correspond to a $O(1)$ bound on the range of public scores, and cannot incorporate enough systematic variation to have multiple tiers.

which any matching protocol that is stable with high probability requires almost as many bits of communication and preference learning (see Theorem 2).

Section 3.1 presents the protocol that achieves this $O^*(\sqrt{n})$ worst-case guarantee. The example in Section 3.2 demonstrates the near optimality of this guarantee.

3.1 Communication-efficient deferred acceptance (CEDA)

The protocol we construct, which achieves the $O^*(\sqrt{n})$ guarantee, is called the *communication-efficient deferred acceptance* (CEDA) protocol. The key idea is to allow workers to better target their applications with the help of signals sent by firms, which help workers identify the firms at which they actually have a non-negligible chance of acceptance. There are two types of signals.

- **Preference signal:** A firm j signals to worker i if it has a high private score ϵ_{ji} for the worker.
- **Qualification requirement signal:** A firm j broadcasts a qualification requirement z_j to the entire market, which specifies the minimum public score a worker who did not receive a preference signal needs to apply to the firm.

The key idea is that if a worker i does not receive a preference signal from a firm and if her public score is below its qualification requirement, then she should not apply to it because she would likely be rejected. We explain this point in more detail after we present the protocol.

Definition 3.1. *A worker i is said to publicly qualify for firm j if her public score meets its qualification requirement, $a_{ji} \geq z_j$. The worker is said to qualify for the firm if she either publicly qualifies for it or receives a preference signal from it.*

The protocol is defined precisely as follows.

Protocol 2. Communication-efficient deferred acceptance (CEDA).

Initialize the qualification requirement for each firm j to the largest integer not exceeding the minimum public score, $z_j = \lfloor a_j \rfloor$.²⁶ There are two phases in CEDA.

²⁶The rounding ensures that the qualification requirement can be communicated with $O(\log n)$ bits.

1. **Preference signaling:** Each firm j places workers according to public scores into unit-ranged bins $[\underline{a}_j, \underline{a}_j + 1)$, $[\underline{a}_j + 1, \underline{a}_j + 2)$, \dots .

The firm sends a preference signal to its top $3\sqrt{n}$ most preferred workers from each bin.²⁷

2. **Deferred acceptance with qualification requirement:** Run the sequential worker-proposing deferred acceptance algorithm (Protocol 1 in Section 2.3) with the following two modifications:

- (a) Workers apply only to firms for which they qualify (see Definition 3.1).
- (b) Each firm j , after every $3e \log(n)\sqrt{n}$ applications from publicly qualified applicants, increases its qualification requirement z_j by 1 and broadcasts this increase to the market by sending a qualification requirement signal.

Theorem 1. *In any separable market, the CEDA protocol (Protocol 2) is a matching protocol that is stable with high probability. Its communication cost and preference learning cost are both at most $O^*(\sqrt{n})$ per agent in the worst case.*

3.1.1 Sketch of the proof

The formal proof of Theorem 1 is in Appendix A. The main reason behind this result is what we call CEDA's *no-false-negatives property*. This property allows CEDA to cut down on unnecessary communication and preference learning costs.

Lemma 3.2 (no-false-negatives property). *With high probability, throughout the running of CEDA, if a worker does not qualify for a firm at a certain time, then she will not be accepted if she had applied to the firm at that time.*

This property implies that with high probability, CEDA prevents only applications that would have been rejected anyway. So the sequence of acceptances in CEDA exactly matches that in the DA protocol, and CEDA succeeds in finding the worker-optimal stable match.

The high-level intuition of why the property holds is as follows. There are two possible reasons for a worker to be desirable to a firm: she has either a high public score or a high

²⁷If there are fewer than $3\sqrt{n}$ workers in the bin, then the firm sends a preference signal to all of them.

private score. The definition of “not qualifying” (inverse of Definition 3.1) rules out both, so workers who do not qualify might as well not apply.

More precisely, define q_j to be the top $\frac{1}{\sqrt{n}}$ th fractile of the private score for firm j , $q_j = F_j^{-1}\left(1 - \frac{1}{\sqrt{n}}\right)$. Define x_j to be firm j 's latent utility for the current tentative match in CEDA. (This is initialized to u_{j0} and increases whenever the firm tentatively accepts a new worker.) We show in Appendix A that CEDA satisfies the following two properties with high probability.

1. Every worker i whose private score for firm j is higher than q_j receives a preference signal.
2. For every firm j , the tentative match value x_j increases sufficiently quickly relative to z_j , so that the following invariant holds throughout the running of the CEDA protocol,

$$z_j \leq \max(\underline{a}_j, x_j - q_j), \quad (4)$$

which implies that if a worker does not publicly qualify for firm j , then her public score satisfies $a_{ji} < x_j - q_j$. This invariant holds due to the following concentration bound: the maximum private score among $3e \log(n)\sqrt{n}$ independent draws of F_j is at least $q_j + 1$ with high probability.

The two points above together imply the no-false-negatives property, because when worker i does not qualify for firm j , then the first property implies that $\epsilon_{ji} < q_j$, and the second implies that $a_{ji} < x_j - q_j$, so the firm's utility for the worker satisfies

$$u_{ji} = a_{ji} + \epsilon_{ji} < x_j,$$

and thus the worker would not have been accepted anyway.

Having explained the intuition as to why CEDA successfully computes the worker-optimal stable match with high probability, we now explain the reasoning behind the $O^*(\sqrt{n})$ bounds in communication and preference learning cost. This follows from the observation that in CEDA, both communication and preference learning can be upper-bounded by the number of signals sent and the number of applications. By construction, each firm sends only $O^*(\sqrt{n})$ preference signals and $O^*(1)$ qualification requirement signals. This implies that it can receive only $O^*(\sqrt{n})$ applications from workers who do not publicly qualify, and $O^*(\sqrt{n})$ applications from workers who do publicly qualify. This last point follows from the observation that there

are at most $O^*(1)$ updates to the qualification requirement for each firm by Assumption 2.4, and for each update, there are at most $O^*(\sqrt{n})$ applications from publicly qualified workers. The argument is explained in more detail in Appendix A.

3.2 The optimality of the $O^*(\sqrt{n})$ guarantee

The following theorem shows that the $O^*(\sqrt{n})$ guarantee of CEDA is asymptotically near optimal.

Theorem 2. *There exists a separable market for which any matching protocol that is stable with high probability requires a communication cost of $\Omega(\sqrt{n})$ per agent and a preference learning cost of $\Omega(\sqrt{n}/\log n)$ per agent.*²⁸

The formal proof of this is in Appendix B, but we explain the main intuitions here. Consider the following example of a separable market: there are n workers and n firms, which are all ex ante identical. The preferences of both workers and firms follow a separable structure, with all public scores being 0, and all private scores being drawn from an exponential distribution with rate parameter 1. Each agent has an outside option of value $\frac{\log n}{2}$, which implies that each agent finds a uniformly random subset of about \sqrt{n} partners acceptable.

For this separable market, we show in Appendix B that any matching protocol that is stable with high probability must use at least $\Omega(n)$ bits of communication per agent. The main ideas are as follows. First, we show that any such matching protocol must approximately identify if each member of worker-firm pair finds the other mutually acceptable. (The precise definition of “approximately identify” is technical and is presented in the formal proof.) The reason is that there are $O(n)$ such mutually acceptable pairs, and a significant fraction of them must be matched in all stable matchings, since $O(n)$ agents are matched in total. There are n^2 worker-firm pairs in total. For each pair, we use ideas from information complexity theory (see Braverman (2015)) to show that the protocol must use $\Omega(1/\sqrt{n})$ bits of communication on average to approximately determine mutual acceptability. This implies that $n^2 \times \Omega(1/\sqrt{n}) = \Omega(n^{\frac{3}{2}})$ total bits of communication are needed, which implies the $\Omega(\sqrt{n})$ per agent bound.

Finally, we show that this $\Omega(\sqrt{n})$ lower bound on communication cost immediately implies a $\Omega(\sqrt{n}/\log n)$ lower bound on preference learning cost. The reason is that any protocol

²⁸Our lower bound allows the communication protocol to be adaptive, to be randomized, and to broadcast messages to every agent.

that uses Q choice function queries can be modified into a communication protocol with a communication cost of $O(Q \log n)$ bits.

4 Two-round protocol with private communication

The CEDA protocol in Section 3.1 is sequential: a worker’s decision about what firms to apply to depends on the firms’ current qualification requirements, which in turn depend on other workers’ application decisions. Implementing such a protocol may create undesirable congestion, as agents need to wait for other agents to act before knowing what preference information to learn next and how to act next. In this section, we explore the possibility of *simultaneous* protocols, in which the dependence of each agent’s action on prior actions is minimized. In particular, we consider *two-round protocols*.²⁹ In the first round, agents simultaneously signal to various partners, and in the second round everyone reports a partial preference list to a central matchmaker, based on signals received during the first round. One interpretation of the central matchmaker is a centralized clearinghouse, as in the National Residency Matching Program (NRMP). Another interpretation is that it is a proxy for a decentralized matching process, after which agents arrive at a matching in which no pair of agents who contacted one another form a blocking pair.

It turns out that one can construct a separable market in which any two-round protocol of this form that is stable with high probability requires $\Omega(n)$ bits of communication per agent (see Appendix G). As a result, we introduce a simpler model of matching markets, which we call *tiered random markets*. In such markets, agents are partitioned into tiers, and each agent prefers better tiers to worse tiers and has uniformly random preferences among agents in a given tier (see Section 4.1). This model still allows both vertical and horizontal differentiation, and it yields clean insights about what kind of signals are most informative. We present our two-round protocol in Section 4.4.

This protocol has an additional advantage: it *uses only private messages*, which are messages visible only to a sender and a receiver, but not to anyone else. The formal definition is as follows. Requiring messages to be private models markets in which there is no efficient way to broadcast a particular message to all agents simultaneously.

Definition 4.1. *A communication protocol Π is said to use only private messages if*

²⁹We need at least two rounds for agents to be able to respond to signals from others.

every message specifies the identity of a receiving agent and is visible only to that agent. Furthermore, each message an agent sends can depend only on the history of messages that the agent has seen, but not on messages the agent has not seen.

As in the definition of a communication protocol in Section 2.3, the protocol still observes all messages and chooses when to terminate with an output and who should send the next message.³⁰

4.1 Tiered random markets

A *tiered market* is a two-sided matching market $(I, J, \omega, K, \mathcal{P})$ in which agents are partitioned into commonly known tiers, and every agent prefers partners from a better tier to those from a worse tier, and has uniformly random preferences for partners within a given tier.³¹

Precisely speaking, there are $K \geq 1$ tiers of workers and $L \geq 1$ tiers of firms. The state of the world $\omega = (s, t)$, where s is a K -dimensional vector of positive integers, and t is an L -dimensional vector of positive integers, satisfying

$$0 =: s_0 < s_1 < \cdots < s_K = n_I,$$

and $0 =: t_0 < t_1 < \cdots < t_L = n_J.$

Let $I_k = \{s_{k-1} + 1, s_{k-1} + 2, \dots, s_k\}$. This denotes the k th tier of workers. I_1 is the best tier, and I_K the worst. Given worker $i \in I$, let $k(i)$ denote the tier of the worker. This is the unique k such that $s_{k-1} < i \leq s_k$. Similarly, let $J_l = \{n_I + t_{l-1} + 1, \dots, n_I + t_l\}$ denote the l th tier of firms. (We add n_I because we index the firms from $n_I + 1$ to $n_I + n_J$.) For every firm j , let $l(j)$ denote the tier of the firm. The a priori knowledge K_a of every agent $a \in I \cup J$ is $\{\omega\}$. The true preference distribution $\mathcal{P}(\omega)$ is such that the preference realization R_i of each worker $i \in I$ is a uniformly random permutation of J_1 , followed by a uniformly random permutation of J_2 , and so on. We assume that every firm is acceptable to the worker. The preferences of firms are defined analogously.

Tiered random markets model markets in which the vertical differentiation is coarse and the horizontal differentiation is idiosyncratic. For example, in the academic job market for

³⁰One can interpret this as the existence of an additional “operator” through which all messages are routed, so the operator sees all messages and can choose who should send the next message.

³¹This preference structure is called “block-correlated” in Coles et al. (2013).

certain subfields, one may argue that departments are clustered into quality-differentiated tiers, and every applicant prefers better tiers but preferences within each tier are driven by personal preferences that can be modeled as essentially random. Furthermore, applicants may also be clustered into tiers based on publication record and school of origin, with different departments having essentially random preferences within each tier based on their particular needs at the moment.

4.2 Better and worse positions

We define a notion of relative competitive position, which intuitively indicates who has stronger market power in a tiered random market. We say that a worker i is in a *weakly worse position* than firm j if $s_{k(i)} \geq t_{l(j)}$. In other words, there are weakly more workers in tiers as good as i than there are firms in tiers as good as j . Similarly define this for firms.

These definitions are motivated by the following result from previous work. In a *uniformly random market*, which is a special case of tiered random markets with one tier on each side, the DA algorithm terminates quickly when the proposing side has weakly fewer agents, but not when it has strictly more agents. (For example, Ashlagi et al. (2017) show that in a market with $n - 1$ workers and n firms, the average number of applications in the worker-proposing DA algorithm is about $\log n$ per agent, whereas the average number of applications in the firm-proposing DA is about $\frac{n}{\log n}$ per agent.) This implies that it is more efficient in terms of communication to have the side with fewer agents do the proposing.

The protocol we propose will have only agents initializing contact with partners in worse positions than themselves.

4.3 Example and intuition

Before giving the full protocol, we consider a simple numerical example, which will illustrate the main insights. Suppose that there are two tiers of workers of 50 workers each. We call these the top workers and bottom workers respectively. Similarly, there are two tiers of firms, which we call the top and bottom firms. There are 20 top firms and 90 bottom firms, as illustrated in Figure 1.

The first observation is that the tier structure precludes certain matches in a stable matching. For example, in every stable matching, every top firm must be matched with a top worker. The reason is that there are more top workers than top firms, and any top

worker who is not matched with a top firm would like to be. Thus, the bottom workers in this example have no chance whatsoever of being matched with a top firm. More generally, in a tiered random market, worker i can be matched with firm j only if their tiers overlap: $s_{k(i)-1} < t_{l(j)}$ and $s_{k(i)} > t_{l(j)-1}$.

The second observation is that it is more efficient to have the agents signal partners of worse positions than themselves. In the example, it is more efficient for the top firms to signal to the top workers than the reverse, because there are fewer top firms than top workers. Similarly, it is more efficient for bottom firms to wait for signals from workers, because when we take out the top workers who will be matched to the top firms, there are $100 - 20 = 80$ workers left, and 90 bottom firms. So the previous results for uniformly random markets suggest that it is better for the workers to signal. These directions of signaling are shown in the arrows in Figure 1.

The third observation is that because preference signals are sent in parallel, certain agents may have to send extra signals to account for the fact that potential partners may already be taken up by competitors from better tiers. For example, for each of the bottom workers, there are 90 bottom firms she can signal, but 30 of these will end up matching with a top worker, against whom she has no chance. So for every three signals she sends, one would be essentially wasted. This means that she should amplify the number of signals she sends by a factor of $\frac{3}{2}$.

For certain agents, this amplification effect may be large. For example, consider the case in which there is a single tier of n firms and there are n tiers of one worker each, so that the workers are completely vertically differentiated. In this example, the very bottom worker may need to signal up to order n firms, because of the amplification effect. However, one can show that the *average* amplification needed is only $O(\log n)$ in this case.³² For arbitrary tier structures, one can show that the average amplification needed is always small.

4.4 The targeted signaling protocol

In this section, we describe a two-round protocol for tiered random markets that is stable with high probability and that uses only private messages. We begin with a high-level summary. The protocol designates for each agent a *target tier* based on the tier structure alone (Definition 4.2). In the first round, every agent signals a certain number of favorite

³²The reason is that the worker ranked number k needs to send only $O(\frac{n}{n-k+1})$ signals.

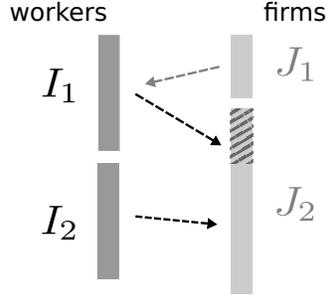


Figure 1: An example of a tiered random market (see Section 4.3). There are two tiers of workers, I_1 and I_2 , and two tiers of firms, J_1 and J_2 . The height of each rectangle corresponds to the number of agents in that tier. The arrows show the direction of the most informative signaling. The top workers I_1 should signal to the bottom firms J_2 , and wait for the top firms to signal to them. The bottom workers I_2 should signal to the bottom firms as well, but they need to amplify the number of signals they send because some of the bottom firms would already be taken by top workers.

partners within the agent’s target tier. A signal intuitively represents initiating contact in a decentralized job market. In the second round, agents submit a partial preference ranking of partners among those they signaled or signaled to them, and the protocol runs the DA algorithm using the partial preferences collected. (The DA algorithm is for concreteness and can be interpreted as a proxy for a decentralized matching process.)

Definition 4.2. *The target tier of a worker i is the best tier of firms that is in a weakly worse position than she is. In other words, this is J_l , where*

$$l = \min\{l : t_l \geq s_{k(i)}\}.$$

Similarly define the target tier of each firm j .

In the example in Figure 1, the target tiers are indicated by the arrows, so that the target tier of top firms is top workers, the target tier of top workers is bottom firms, and so on. The bottom firms do not have a target tier because they are in the worst position possible. As explained in Section 4.3, it is most efficient to limit signals to one’s target tier. However, the number of signals sent needs to be amplified inversely proportional to the fraction of agents left over from competitors in better tiers. This is precisely stated below.

Definition 4.3. For each worker i in tier k with target tier l , define her relative competitiveness for her target tier

$$\rho(i) = \min \left(1, \frac{t_l - s_{k-1}}{t_l - t_{l-1}} \right).$$

This is the proportion of the target tier that will not be matched to workers from better tiers. Similarly define $\rho(j)$ for firms.

Definition 4.4. Define the target number of an agent $a \in I \cup J$ as

$$r(a) = \frac{24 \log^2 n}{\rho(a)}.$$

In the example, the relative competitiveness of top firms and top workers is 1, whereas the relative competitiveness of the bottom workers is the proportion of bottom firms that are unshaded in Figure 1, which is equal to $\frac{60}{90} = \frac{2}{3}$. The most important feature of the target number is that it is inversely proportional to $\rho(a)$, which matches the level of amplification in the example. The $O(\log^2 n)$ multiplier accounts for competition from the agent's own tier.

The protocol is precisely defined as follows.

Protocol 3. Targeted signaling protocol.

The protocol has two rounds and uses only private messages.

1. **Signaling round:** every agent $a \in I \cup J$ signals to the agent's favorite $r(a)$ partners in the agent's target tier.
2. **Matching round:** every agent submits to the protocol a partial preference ranking of partners, with the ranking restricted to partners who either signaled to the agent or to whom the agent signaled. The protocol then runs the worker-proposing DA algorithm based on the partial rankings, and outputs the resultant matching.³³

Theorem 3. For any tiered random market, the targeted signaling protocol is a matching protocol that uses only private messages and is stable with high probability. Its average communication cost is $O(\log^4 n)$ per agent, and its average preference learning cost is $O(\log^3 n)$ per agent.

³³The worker-proposing DA algorithm is run off-line, using only the partial preferences collected already, and does not require further communication with agents. The firm-proposing DA can also be used.

The proof of Theorem 3 is in Appendix D. The main steps are as follows. First, we show that, with high probability, a certain subset of signals sent in the signaling round contains a stable matching. This result, stated in Lemma D.3, is obtained by proving a new bound on the average rank of agents in any stable matching in unbalanced uniform random markets.³⁴ Second, we show that whenever this subset of signals contains a stable matching, running DA on the partial preferences as in the matching round returns a stable matching. This result uses the structure of tiered markets and the definition of target tier, and is not true in general; even if a set of partial preferences contains a stable matching with respect to full preferences, running DA on these partial preferences may result in a matching that is stable only with respect to the partial preferences. Third, we count the total number of signals and show that it is no more than $O(\log^3 n)$ per agent.

Remark 4.5. *An immediate corollary of Theorem 3 is that in the targeted signaling protocol, with high probability, the number of agents who experience more than \sqrt{n} bits of communication is at most $O^*(\sqrt{n})$.³⁵ Intuitively speaking, some agents may experience a lot of communication, but the number of such agents is a vanishing fraction of the population. In Theorem 8 of Appendix H, we show that having some agents experience $\Omega(\sqrt{n})$ bits of communication is unavoidable by any protocol that uses private messages and is stable with high probability.*

Remark 4.6. *We also have a variant of the targeted signaling protocol that saves a $O(\log n)$ factor: its communication cost is $O(\log^3 n)$ bits per agent and its preference learning cost is $\Theta(\log^2 n)$ per agent. It involves a more complicated formula for the target number than in Definition 4.4. In the interest of a clean exposition, we show the simpler version here.*

Remark 4.7. *The targeted signaling protocol presented above achieves near optimal communication cost. In Theorem 8 of Appendix H, we give an example of a tiered random market in which any matching protocol that uses only private communication and is stable with high probability must use at least $\Omega(\log^2 n)$ bits of communication per agent. This lower bound includes protocols that allow an arbitrary number of rounds of communication. The targeted signaling protocol achieves near optimal cost using only two rounds.*

³⁴The bound, stated in Proposition E.1, is analogous to Theorem 1(i) of Ashlagi et al. (2017), except that the constant is worse but the notion of “high probability” is stronger. The proof of Proposition E.1 is based on the integral formula technique of Pittel (1992) and is much cleaner than the analysis in Ashlagi et al. (2017).

³⁵This follows because the total amount of communication is $O(n \log^3 n) = O^*(n)$.

5 Incentive compatibility

The mechanism induced by the deferred acceptance (DA), even in the setting with full preference elicitation, is not completely incentive compatible: an agent on the non-proposing side may profitably deviate from truthful reporting by truncating his or her preference ranking.³⁶ Nevertheless, the DA mechanism is known to be strategyproof for the proposing side.³⁷ Moreover, assuming truthful reporting by other agents, an agent on the non-proposing side cannot unilaterally deviate (misreport her preferences) and be matched with someone better than her best stable partner.³⁸

We show that, with high probability, all the matching protocols proposed in this paper are in a certain sense “as incentive compatible” as DA with full preference elicitation.

Definition 5.1. *For a given two-sided matching market \mathcal{M} , a matching protocol Π is as incentive compatible as DA with high probability if there exists a function $\delta(n)$ with $\delta(n) \rightarrow 0$ as $n \rightarrow \infty$, such that for any fixed agent, the probability that the agent can unilaterally deviate from the protocol and be matched with someone better than the agent’s best stable partner (assuming true underlying preferences) is at most $\delta(n)$. The probabilities are defined with respect to the randomness of preferences as well as the randomness of the protocol.*

Theorem 4. *For separable markets, the CEDA protocol (Protocol 2) is as incentive compatible as DA with high probability. For tiered random markets, the targeted signaling protocol (Protocol 3) is as incentive compatible as DA with high probability.*

The “with high probability” caveat is needed because, with a small probability, the protocols may fail to find a stable matching. The proof of Theorem 4 is in Appendix F. The proof is based on applying a “Blocking Lemma” by Gale and Sotomayor (1985) to a market in which preferences are restricted to the *subgraph of signals*, which is the set of worker-firm pairs in which at least one member of each pair sent a signal to the other (i.e., a preference signal or an application in the CEDA protocol, or a signal in the targeted signaling protocol). The same lemma is used to prove the incentive properties of the original DA algorithm.³⁹

³⁶See Chapter 4 of Roth and Sotomayor (1990) for a thorough exposition of all these incentive properties of DA.

³⁷See Roth (1982) and Dubins and Freedman (1981).

³⁸I.e., their most preferred partner in all stable matchings. For a proof, see Theorem 4.11 in Roth and Sotomayor (1990) which is originally credited to Demange et al. (1987).

³⁹See pages 92–93 of Roth and Sotomayor (1990).

The new ingredient here is to use properties of the subgraph of signals in our protocols. In particular, we make use of the fact that a single deviating agent has little control over the edges of the subgraph between other agents.

Since both protocols we propose are based on the worker-proposing DA, they are also, with high probability, strategyproof for workers.

Corollary 5.2. *For separable markets, the CEDA protocol is strategyproof for all workers, with high probability. For tiered random markets, the targeted signaling protocol is strategyproof for all workers, with high probability. (The probabilities are defined with respect to the randomness of preferences.)*

Theorem 4 also implies that whenever an agent has a unique stable partner, the probability that the agent can profitably deviate from the protocol vanishes in large markets. Thus, if for each agent the probability that the agent has multiple stable partners is small, then following the protocol is an ϵ -Bayes Nash equilibrium. In tiered random markets, we can prove this condition for “typical” markets when the size of each tier is large, and hence deduce approximate incentive compatibility for the targeted signaling protocol. This is made precise as follows.

Definition 5.3. *A tiered market satisfies general imbalance if the sets $\{t_1, t_2, \dots\}$ and $\{s_1, s_2, \dots\}$ are disjoint.*

Theorem 5. *There exists a function $\delta : \mathbb{N} \rightarrow \mathbb{R}$ satisfying $\delta(y) \rightarrow 0$ as $y \rightarrow \infty$ such that the following holds. For any $y \in \mathbb{N}$, consider any tiered random market satisfying general imbalance in which the number of agents in each tier is at least y . Then for each agent, the probability that the agent can profitably deviate from the targeted signaling protocol is at most $\delta(y)$.*

The proof of this is in Appendix F.

6 Discussion

While we study particular mathematical models in this paper, the general message is broader: a stable matching can be achieved in large markets by informative signaling, despite real-

world limitations in communication.⁴⁰ This supports the use of stable matching as an appropriate equilibrium concept for empirical estimation. Another message is that it is useful to study signaling in economic models not only from a strategic perspective, but also from an information theory perspective, as this gives insights into how one may directly reduce information friction by using more informative signals.

The types of informative signals we uncover for matching markets include the following: signaling partners for whom one has a particular idiosyncratic preference, broadcasting to the market a requirement on observable characteristics, and focusing on initiating contact with easy-to-get partners while waiting for hard-to-get partners to initiate contact. These types of signaling are already present in many real matching markets, and our results highlight their importance to market efficiency.

Our results also give prescriptive guidelines for the design of real matching markets. One application is identifying what information to show and what interactions to allow in online matching markets for labor and dating. For example, the dating platform Tinder shows each user a series of profiles of potential partners and allows users to “like” or “dislike” each profile shown. When two users “like” each other, the platform recommends a match. Previous studies have demonstrated that on this platform, many users can spend a long time indicating interest to many partners, but get no response. A question is how should the platform decide which profiles to show each user, in order to maximize the overall goodness of match while reducing search friction. One thing the platform already does is to estimate a quality rating for each user’s profile, as well as potential goodness of fit. Our results suggest that rather than simply showing users the best profiles for them, it is more efficient to show the best profiles that are “easy-to-get” for them, plus the best profiles that are “hard-to-get” but who previously liked them already.⁴¹

Another application lies in decreasing the logistic overhead of the pre-matching phase in centralized matching markets. In the NRMP for medical fellows, one current problem is that it is costly and time consuming to screen applications and to hold on-site interviews, as doctors in teaching hospitals are already extremely busy.⁴² Our results suggest that one may

⁴⁰See Hitsch et al. (2010), which find that predicted (stable) matches are similar to actual matches in an online dating website.

⁴¹For online platforms in particular, there are other potential ways to reduce communication overhead, such as restricting communication, imposing costs on communication, or charging prices for obtaining information. We do not study these policy levers directly in this paper but they are interesting directions for future work.

⁴²This problem was a topic of discussion at the 2017 annual meeting of the Society of Surgical Oncology,

be able to improve the efficiency of search by having the matching platform suggest interviews based on the likelihood that a particular interview would lead to a stable matching. This can be done by estimating the relative competitiveness of agents based on their profiles and by allowing agents to signal beforehand their preferences for particular partners or types of partners. The details behind such a system require further research, but our current results suggest that something of this nature may be possible.

This paper raises several questions for future research. One question is how many rounds of communications to allow, while keeping the overall level of communication low.⁴³ In Appendix G, we give a simple example of a separable market in which no two-round protocol can reach a stable matching with low communication cost. In that example, the problem would be solved by having an additional, aftermarket round.⁴⁴

We expect that our results extend to the case, in which agents are able to only approximately learn their own preferences. Suppose each agent can only query an *approximate* choice function, which for each subset of partners returns a partner for whom her cardinal utility is within ϵ of her utility to her favorite partner in the subset. Then a suitable version of our communication protocol CEDA should provably return an approximately stable matching, in which at least one member of each blocking pair would improve her utility by no more than ϵ .

Another question is, How heterogeneous do communication and preference learning costs need to be across agents? In the targeted-signaling protocol, while most agents incur a low communication cost (polylogarithmic in the market size), some agents may require a lot of communication. To see this, consider the following example: there are n tiers of one worker each, and one tier of n firms. In this example, the top workers can essentially choose whichever firm they like, but the bottom workers can be matched only with leftover firms after workers from higher tiers have already made their choices. In the targeted-signaling protocol, the bottom worker in this example incurs a linear communication cost. Moreover, we show in Appendix H that for this market, any protocol⁴⁵ that uses only private messages requires some agent to incur a communication cost on the order of the square root of the market size. It would be interesting to systematically study this relationship between

which one of us attended.

⁴³One paper that studies the number of rounds of communication needed in a matching model is Dobzinski et al. (2014). However, they do not consider the stability of the match.

⁴⁴Such an aftermarket is implemented in the NRMP after the main match takes place.

⁴⁵The protocol can have arbitrarily many rounds.

congestion and the market position of particular agents, both theoretically and empirically.

References

- Atila Abdulkadirođlu, Parag A. Pathak, and Alvin E. Roth. The new york city high school match. *American Economic Review*, 95(2):364–367, 2005a.
- Atila Abdulkadirođlu, Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez. The boston public school match. *American Economic Review*, 95(2):368–371, 2005b.
- Atila Abdulkadirođlu, Yeon-Koo Che, and Yosuke Yasuda. Expanding “choice” in school choice. *American Economic Journal: Microeconomics*, 7(1):1–42, 2015.
- A. Abdulkadirođlu, P. A. Pathak, and A. E. Roth. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the nyc high school match. *American Economic Review*, 99(5):1954–1978, 2009.
- Heiner Ackermann, Paul W. Goldberg, Vahab S. Mirrokni, Heiko Röglin, and Berthold Vöcking. Uncoordinated two-sided matching markets. *SIAM Journal on Computing*, 40(1):92–106, 2011.
- Hiroyuki Adachi. A search model of two-sided matching under nontransferable utility. *Journal of Economic Theory*, 113(2):182–198, 2003.
- Nikhil Agarwal. An empirical model of the medical match. *The American Economic Review*, 105(7):1939–1978, 2015.
- Nick Arnosti, Ramesh Johari, and Yash Kanoria. Managing congestion in decentralized matching markets. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, pages 451–451. ACM, 2014.
- Itai Ashlagi, Yash Kanoria, and Jacob D. Leshno. Unbalanced random matching markets: The stark effect of competition. *Journal of Political Economy*, 125(1):69–98, 2017.
- Christopher Avery, Christine Jolls, Richard A. Posner, and Alvin E. Roth. The market for federal judicial law clerks. *The University of Chicago Law Review*, 68(3):793–902, 2001.

- Laszlo Babai, Peter Frankl, and Janos Simon. Complexity classes in communication complexity theory. In *Proceedings of the 27th Annual Symposium on Foundations of Computer Science*, pages 337–347. IEEE, 1986.
- Mark Braverman. Interactive information complexity. *SIAM Journal on Computing*, 44(6): 1698–1739, 2015.
- Mark Braverman, Ankit Garg, Denis Pankratov, and Omri Weinstein. From information to exact communication. In *Proceedings of the 45th Annual ACM Symposium on Theory of Computing*, pages 151–160. ACM, 2013.
- Hector Chade, Jan Eeckhout, and Lones Smith. Search and matching models in microeconomics. Technical report, working paper, 42., 2014.
- Pierre-André Chiappori and Bernard Salanié. The econometrics of matching models. *Journal of Economic Literature*, 54(3):832–861, 2016.
- Eugene Choo and Aloysius Siow. Who marries whom and why. *Journal of Political Economy*, 114(1):175–201, 2006.
- Jen-Hou Chou and Chi-Jen Lu. Communication requirements for stable marriages. In *Proceedings of the 7th International Conference on Algorithms and Complexity*, pages 371–382. Springer, 2010.
- Peter Coles, John Cawley, Phillip B. Levine, Muriel Niederle, Alvin E. Roth, and John J. Siegfried. The job market for new economists: A market design perspective. *The Journal of Economic Perspectives*, 24(4):187–206, 2010.
- Peter Coles, Alexey Kushnir, and Muriel Niederle. Preference signaling in matching markets. *American Economic Journal: Microeconomics*, 5(2):99–134, 2013.
- Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. John Wiley & Sons, 2012.
- John K Dagsvik. Aggregation in matching markets. *International Economic Review*, 41(1): 27–58, 2000.

- Sanmay Das and Emir Kamenica. Two-sided bandits and the dating market. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence*, volume 5, page 19, 2005.
- Gabrielle Demange, David Gale, and Marilda Sotomayor. A further note on the stable matching problem. *Discrete Applied Mathematics*, 16(3):217–222, 1987.
- Shahar Dobzinski, Noam Nisan, and Sigal Oren. Economic efficiency requires interaction. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 233–242. ACM, 2014.
- Lester E. Dubins and David A. Freedman. Machiavelli and the gale-shapley algorithm. *The American Mathematical Monthly*, 88(7):485–494, 1981.
- Federico Echenique and Leeat Yariv. An experimental study of decentralized matching. Working paper, 2013.
- Federico Echenique, Alistair J. Wilson, and Leeat Yariv. Clearinghouses for two-sided matching: An experimental study. *Quantitative Economics*, 7(2):449–82, 2016.
- Jan Eeckhout. Bilateral search and vertical heterogeneity. *International Economic Review*, 40(4):869–887, 1999.
- Andrey Fradkin. Search, matching, and the role of digital marketplace design in enabling trade: Evidence from airbnb. Working paper, 2017.
- David Gale and Lloyd L. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
- David Gale and Marilda Sotomayor. Some remarks on the stable matching problem. *Discrete Applied Mathematics*, 11(3):223–232, 1985.
- Alfred Galichon and Bernard Salanié. Cupids invisible hand: Social surplus and identification in matching models. Working paper, 2015.
- Alfred Galichon, Scott D. Kominers, and Simon Weber. Costly concessions: An empirical framework for matching with imperfectly transferable utility. Working paper, 2016.

- Yannai A. Gonczarowski, Noam Nisan, Rafail Ostrovsky, and Will Rosenbaum. A stable marriage requires communication. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1003–1017. SIAM, 2015.
- Guillaume Haeringer and Myrna Wooders. Decentralized job matching. *International Journal of Game Theory*, 40(1):1–28, 2011.
- John W. Hatfield and Paul R. Milgrom. Matching with contracts. *The American Economic Review*, 95(4):913–935, 2005.
- Friedrich A. Hayek. The use of knowledge in society. *The American Economic Review*, 34(4):519–530, 1945.
- Günter J Hitsch, Ali Hortaçsu, and Dan Ariely. Matching and sorting in online dating. *The American Economic Review*, 100(1):130–163, 2010.
- John J. Horton. The effects of algorithmic labor market recommendations: Evidence from a field experiment. *Journal of Labor Economics*, 35(2):345–385, 2017.
- Leonid Hurwicz. The design of mechanisms for resource allocation. *The American Economic Review*, 63(2):1–30, 1973.
- Nicole Immorlica and Mohammad Mahdian. Marriage, honesty, and stability. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 53–62, 2005.
- John H. Kagel and Alvin E. Roth. The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment. *The Quarterly Journal of Economics*, 115(1):201–235, 2000.
- Bala Kalyanasundaram and Georg Schnitger. The probabilistic communication complexity of set intersection. *SIAM Journal on Discrete Mathematics*, 5(4):545–557, 1992.
- Yash Kanoria and Daniela Saban. Facilitating the search for partners on matching platforms: Restricting agent actions. Working paper, 2017.
- Alexander S. Kelso and Vincent P. Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society*, 50(6):1483–1504, 1982.

- Donald E. Knuth. *Mariages stables et leurs relations avec d'autres problèmes combinatoires: introduction à l'analyse mathématique des algorithmes*. Montréal: Presses de l'Université de Montréal, 1976.
- Fuhito Kojima and Parag A. Pathak. Incentives and stability in large two-sided matching markets. *American Economic Review*, 99:608–627, 2009.
- Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 2006.
- Stephan Lauermaun and Georg Nöldeke. Stable marriages and search frictions. *Journal of Economic Theory*, 151:163–195, 2014.
- Robin S. Lee and Michael Schwarz. Interviewing in two-sided matching markets. Technical report, National Bureau of Economic Research, 2017.
- SangMok Lee. Incentive compatibility of large centralized matching markets. *The Review of Economic Studies*, 84(1):444–463, 2017.
- Soohyung Lee and Muriel Niederle. Propose with a rose? signaling in internet dating markets. *Experimental Economics*, 18(4):731–755, 2015.
- D. G. McVitie and L. B. Wilson. The stable marriage problem. *Communications of the*, 14(7):486–490, July 1971.
- Konrad Menzel. Large matching markets as two-sided demand systems. *Econometrica*, 83(3):897–941, 2015.
- Kenneth Mount and Stanley Reiter. The informational size of message spaces. *Journal of Economic Theory*, 8(2):161–192, 1974.
- Muriel Niederle and Leeat Yariv. Decentralized matching with aligned preferences. Technical report, National Bureau of Economic Research, 2009.
- Joana Pais. Incentives in decentralized random matching markets. *Games and Economic Behavior*, 64(2):632–649, 2008.
- Marcin Peski. Large roommate problem with non-transferable random utility. working paper, 2015.

- Boris Pittel. The average number of stable matchings. *SIAM Journal on Discrete Mathematics*, 2(4):530–549, 1989.
- Boris Pittel. On likely solutions of a stable marriage problem. *The Annals of Applied Probability*, 2(2):358–401, 1992.
- Alexander Razborov. On the distributed complexity of disjointness. *TCS: Theoretical Computer Science*, 106:385–390, 1992.
- Alvin E. Roth. The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7(4):617–628, 1982.
- Alvin E. Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of political Economy*, 92(6):991–1016, 1984.
- Alvin E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, 2002.
- Alvin E. Roth. What have we learned from market design? *The Economic Journal*, 118(527):285–310, 2008.
- Alvin E. Roth and Elliott Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *The American Economic Review*, 89(4):748, 1999.
- Alvin E. Roth and Oliveira Sotomayor. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, 1990.
- Alvin E. Roth and John H. Vande Vate. Random paths to stability in two-sided matching. *Econometrica: Journal of the Econometric Society*, 58(6):1475–1480, 1990.
- Alvin E. Roth and Xiaolin Xing. Jumping the gun: Imperfections and institutions related to the timing of market transactions. *The American Economic Review*, 84(4):992–1044, 1994.
- Ilya Segal. The communication requirements of social choice rules and supporting budget sets. *Journal of Economic Theory*, 136(1):341–378, 2007.

Morten Sørensen. How smart is smart money? a two-sided matching model of venture capital. *The Journal of Finance*, 62(6):2725–2762, 2007.

A Proof of correctness and efficiency of CEDA (Theorem 1)

In this appendix, we prove Theorem 1 from Section 3.1, which states that in any separable market, the CEDA protocol is a matching protocol that is stable with high probability (see Definition 2.1), and that it uses only communication and preference learning costs of $O^*(\sqrt{n})$ per agent. After proving this, we give a generalization of CEDA that works when the range of public scores is upper bounded by a general function $g(n)$, instead of the $O^*(1)$ upper bound assumed in Section 2.4.

Proof of Theorem 1. We first prove that CEDA is a stable with high probability in separable markets. (In other words, it succeeds in finding a stable matching with high probability.) Define q_j to be the top $\frac{1}{\sqrt{n}}$ -th fractile of private score for firm j , $q_j = F_j^{-1}(1 - \frac{1}{\sqrt{n}})$. Define x_j to be firm j 's latent utility for the current tentative match. (This is initialized to u_{j0} and increases with every tentative acceptance by the firm.) The proof relies on the two following lemmas, which we will prove later.

Lemma A.1. *With probability at least $1 - o(\frac{1}{n})$, for every firm j , every worker i whose private score for the firm is at least q_j receives a preference signal from the firm.*

Lemma A.2. *With probability at least $1 - o(1/n)$, throughout the running of the CEDA protocol, we have the following invariant for every firm j ,*

$$z_j \leq \max(\underline{a}_j, x_j - q_j) \tag{5}$$

Given Lemma A.1 and A.2, we show that CEDA successfully computes the worker-optimal stable match with high probability. The argument is by showing what is referred to as the *no-false-negatives property* in Section 3.1, which is that with high probability, CEDA never prevents an application that would have been accepted. The usefulness of this property comes from the fact that when CEDA does not prevent any application that would have been

accepted, then CEDA reproduces the outcome of the standard DA algorithm (Protocol 1), which is the worker-optimal stable match.

The no-false-negatives property follows from the two lemmas above. This is because the only applications CEDA prevents are from worker-firm pairs (i, j) in which worker i does not qualify to firm j . By definition, this means that the worker does not publicly qualify for the firm, and has not received a preference signal from the firm. Suppose that we are in the $1 - o(\frac{1}{n})$ fraction of preferences in which the statements in both Lemma A.1 and A.2 are true. Then the first clause above implies that $a_{ji} < z_j$, which implies that $z_j > \underline{a}_j$, so by Inequality (5), $a_{ji} < z_j < x_j - q_j$. The second clause above implies that $\epsilon_{ji} < q_j$. Together, these two clauses implies that

$$u_{ji} = a_{ji} + \epsilon_{ji} < x_j - q_j + q_j = x_j,$$

so worker i would not have been accepted by firm j even if the worker applied. This implies that with probability $1 - o(\frac{1}{n})$ (with probability defined on the randomness in the preferences), CEDA produces the worker-optimal stable match.

Having established that CEDA is stable with high probability in separable markets, we now prove the $O^*(\sqrt{n})$ bounds on the average communication cost per agent, as well as on the average preference learning cost per agent.

First, we bound the number of signals sent per firm. For preference signals, this is at most $O^*(\sqrt{n})$ because there are $O^*(1)$ bins and $3\sqrt{n}$ signals per bin. For qualification signals, this is at most $O^*(1)$ because this is the maximum range of public scores for any firm and the qualification requirement increases by 1 for every signal. (When the qualification requirement exceeds the maximum public score, then the increases in qualification requirement would also stop because there are no publicly qualified applicants.)

The bounds on the number of signals sent per firm imply a $O^*(\sqrt{n})$ bound on the number of applications received by each firm. This is because the number of applications from workers who do not publicly qualify is upper bounded by the number of preference signals. The number of applications from workers who do publicly qualify is upper bounded by $3e \log n \sqrt{n}$ times the number of qualification requirement updates. Both of these are $O^*(\sqrt{n})$.

To bound preference learning cost, observe that the only choice function queries needed in CEDA are for preference signals, applications, and responses to an application, and each of these requires exactly one choice function query. This proves the $O^*(\sqrt{n})$ per agent bound on preference learning cost. For communication cost, observe that each preference signal,

application, and response to an application can be communicated in $O(\log n)$ bits, as this suffices to encode the identity of an agent. Each qualification requirement signal can also be represented by $O^*(1)$ bits, since each qualification requirement z_j only takes integer values and have a range of $O^*(1)$ bits.

To complete the proof of Theorem 1, we prove Lemmas A.1 and A.2 below. Both of these use Lemma A.3, which we derive first.

Lemma A.3. *If a random variable Z is distributed according to CDF F with bounded hazard rate*

$$h(x) = \frac{F'(x)}{1 - F(x)} \leq 1 \quad \forall x \in \mathbb{R}.$$

Then for all $x \in \mathbb{R}$,

$$\frac{\mathbb{P}(Z \geq x + 1)}{\mathbb{P}(Z \geq x)} \geq \frac{1}{e}.$$

Proof of Lemma A.3. Define $\phi(x) = \log(1 - F(x))$. Note that bounded hazard rate implies that the derivative $\phi'(x) \geq -1$. The result follows from the fact that the desired conditional probability is simply $\exp(\phi(x + 1) - \phi(x))$. \square

Proof of Lemma A.1. For a given firm j , suppose on the contrary that there exists a worker i whose private score exceeds q_j but who does not get a preference signal from the firm. Then there must be at least $3\sqrt{n}$ other workers in the same bin as worker i who do receive a preference signal. This in turn implies that there are $3\sqrt{n}$ other workers with private score at least $q_j - 1$.

Let random variable Y denote the number of workers in the bin with private score of at least $q_j - 1$. Define $y = E[Y]$. By Lemma A.3 (with x in the Lemma being $q_j - 1$), we get $y \leq e\sqrt{n}$. Furthermore, $y \geq \sqrt{n}$ by definition of q_j . Let $M = 3\sqrt{n}$, and $\beta = \frac{3-e}{e}$. Note that $(1 + \beta)y \leq M$. By Chernoff bound,

$$\mathbb{P}(Y \geq M) \leq \mathbb{P}(Y \geq (1 + \beta)y) \leq \exp\left(-\frac{\beta^2}{3}y\right) \leq \exp\left(-\frac{\beta^2}{3}\sqrt{n}\right) = o\left(\frac{1}{n^2}\right). \quad (6)$$

Thus, after a union bound on the $O^*(1)$ bins per firm and $O(n)$ firms, we get that the desired result. \square

Proof of Lemma A.2. Fix firm j , the invariant on z_j (see Inequality (5)) is initially satisfied by definition, as $z_j = \lfloor \underline{a}_j \rfloor \leq \max(\underline{a}_j, x_j - q_j)$. There are at most $O^*(1)$ increases to z_j

throughout the running of CEDA. Since the right hand side of the invariant can only increase, it suffices to upper bound the probability that $z_j > x_j + q_j$ after each of these increases.

Now, suppose that the qualification requirement z_j is equal to y at some point in CEDA. We show that with probability at least $1 - \frac{1}{n^3}$, after $3e \log n \sqrt{n}$ applications from publicly qualified workers, the tentative match value x_j is at least $y + q_j + 1$. This is because if x_j is less than this after so many applications from publicly qualified workers, then it must be that all of those workers had private score less than $q_j + 1$. By Lemma A.3, the chance that any worker has private score at least $q_j + 1$ is at least $\frac{1}{e\sqrt{n}}$. So the chance that not one of $3e \log n \sqrt{n}$ workers had such a high private score is at most

$$\left(1 - \frac{1}{e\sqrt{n}}\right)^{3e \log n \sqrt{n}} \leq \frac{1}{n^3}.$$

The above statement implies that regardless of what the value of $y = z_j$ is at a certain time, after $3e \log n \sqrt{n}$ applications from publicly qualified workers, we have that with probability at least $1 - \frac{1}{n^3}$, the qualification requirement at that time is

$$z_j = y + 1 \leq x_j - q_j \leq \max(\underline{a}_j, x_j - q_j).$$

Using this argument, and counting from the first application from a publicly qualified applicant, we get that with probability at least $1 - \frac{O^*(1)}{n^3} = 1 - o(\frac{1}{n^2})$, the invariant is satisfied after every qualification requirement update. A union bound on the $O(n)$ firms yields the desired result. \square

This completes the proof of Theorem 1. \square

A.1 Extension of CEDA to larger range of public scores

Suppose now that the range in public scores for a firm is upper bounded by a general function $g(n)$, instead of by the $O^*(1)$ bound assumed in Assumption 2.4. Finding a stable match in such a market is a more difficult problem because there is a greater variation of possible unknowns. In fact, if we allow $g(n)$ to be as high as $O^*(n)$, then one can show that even with private scores satisfying bounded hazard rate, it is possible to embed the worst-case examples of Gonczarowski et al. (2015) into our model so that any matching protocol that is stable with high probability must use at least $\Omega(n)$ bits of communication per agent.

Nevertheless, we can generalize CEDA to cases in which $g(n) = o(n)$, and still provide a $O^*(\sqrt{ng(n)})$ guarantee on communication and preference learning costs. This is nontrivial

because $O^*(\sqrt{ng(n)}) = o(n)$, so this bypasses the impossibility result for arbitrary markets (Proposition 2.3). The protocol is as follows.

Protocol 4. Generalized CEDA

Initialize the qualification requirement for each firm j to $z_j = \lfloor \underline{a}_j \rfloor$. There are two phases.

1. **Preference signaling:** *Each firm j bins workers according to public scores into unit-ranged bins $[\underline{a}_j, \underline{a}_j + 1)$, $[\underline{a}_j + 1, \underline{a}_j + 2)$, \dots . For each bin, let the number of workers be l . The firm sends a preference signal to its top*

$$M(n, l) = \max \left((\log n)^2, l \sqrt{\frac{g(n)}{n}} \right) \tag{7}$$

most preferred workers from the bin.

2. **Deferred acceptance with qualification requirement:** *Proceed as in step 2 of the CEDA protocol (see Section 3.1), except that we change the number of publicly qualified applications received before increasing the qualification requirement by one to*

$$3e \log n \sqrt{\frac{n}{g(n)}}.$$

Note that the only changes when generalizing CEDA are the number of preference signals sent in each bin, and the number of applications to wait for before each update of the qualification requirement. The underlying framework is the same.

Theorem 6. *In any separable market in which the range of public scores is at most $g(n)$, the generalized CEDA protocol (Protocol 4) is a matching protocol that is stable with high probability. In the worst case, its communication cost and preference learning cost are $O^*(\sqrt{ng(n)})$ per agent.*

Proof of Theorem 6. Define q_j to be now the top $\sqrt{\frac{g(n)}{n}}$ -th quantile of private scores. By the argument in the proof of Theorem 1, to show that the generalized CEDA protocol is stable with high probability, it suffices to prove that Lemmas A.1 and A.2 hold in this new context with this new definition of q_j .

For Lemma A.1, the same argument works except we need to modify the Chernoff bound (Expression (6)) to handle bins with few workers. As in the previous proof of Lemma A.1, for

a particular bin with l workers, define Y to be the number of workers in the bin with private score at least $q_j - 1$, and $y = E[Y]$. By Lemma A.3, $y \leq e\sqrt{\frac{g(n)}{n}}l$. Define $\beta = \frac{M(n,l)}{\mu} - 1$. By definition of M (Equation 7), $\beta \geq \frac{3-e}{e}$, and $\beta y \geq \frac{3-e}{3}M(n,l) = \Omega(\log^2 n)$. By Chernoff bound,

$$\mathbb{P}(Y \geq M(n,l)) \leq \max\left(e^{-\frac{\beta^2 y}{3}}, e^{-\frac{\beta^2 y}{2+\beta}}\right) \leq \exp\left(-\frac{3-e}{e}\frac{\beta y}{3}\right) \leq \exp(-\Omega(\log^2(n))) = o\left(\frac{1}{n^3}\right).$$

As before, a union bound over all $g(n)$ bins for each firm and the $O(n)$ firms yields the desired result.

For Lemma A.2, the same argument as in the previous proof applies. Concretely speaking, suppose that the tentative match value z_j is equal some value y at some point. With probability at least $1 - \frac{1}{n^3}$, after seeing $3e \log n \sqrt{\frac{n}{g(n)}}$ applications from publicly qualified workers, the tentative match value x_j is at least $y + q_j + 1$. This is because the chance that any worker has a private score of at least $q_j + 1$ is at least $\frac{1}{e}\sqrt{\frac{g(n)}{n}}$, and the chance that none of these publicly qualified applicants have such a high private score is at most

$$\left(1 - \frac{1}{e}\sqrt{\frac{g(n)}{n}}\right)^{3e \log n \sqrt{\frac{n}{g(n)}}} \leq \frac{1}{n^3}.$$

After establishing Lemmas A.1 and A.2 for this new setting with range of public score being $g(n)$, we get from the argument in the proof of Theorem 1 that the generalized CEDA protocol is stable with high probability for such markets.

As before, to prove the bound on communication and preference learning costs, it suffices to bound the number of signals and applications per agent. By definition and by the assumption that $g(n) = o(n)$, each firm sends at most $O^*(\sqrt{ng(n)})$ preference signals. This also bounds the number of applications from workers who are not publicly qualified. Each firm also sends at most $g(n)$ qualification requirement signals, and for each such signal, receives at most $O^*\left(\frac{n}{g(n)}\right)$ applications from publicly qualified workers. So the total number of signals and applications is at most $O^*(\sqrt{ng(n)})$ per agent, which is what we needed to prove. \square

B Proof of Near Optimality of CEDA (Theorem 2)

We obtain the communication lower bound in a model that is even stronger than the broadcast model. Consider the following two-party relaxation of the problem of finding a stable

matching. Alice controls all the workers, and Bob controls all the firms. The workers' and firms' preferences are generated according to the model. Alice and Bob want to figure out a stable matching by communicating with each other. Note that if there is a distributed broadcasting protocol that uses a total of B bits of communication, then Alice and Bob can simulate it using B bits of communication: Alice will simulate all the workers' messages, and Bob will simulate all the firms' messages. Note that the converse is not true, since Alice's messages are allowed to depend on the preferences of all workers simultaneously (which amounts to having "free" communication among workers). We will show an example where $\Omega(n^{3/2})$ communication between Alice and Bob is necessary, which will immediately imply an $\Omega(n^{3/2}/n) = \Omega(\sqrt{n})$ lower bound on average communication per agent needed to solve the original (harder) problem.

In fact, we show our lower bound under a further restriction of the model with the workers' preferences being stochastic and similar to the firms' preferences. The construction is as follows. There are n workers and n firms. Let v_{ij} be worker i 's latent utility for firm j and u_{ji} be firm j 's latent utility for worker i . Let both be distributed independently according to $Exp(1)$, which is the exponential distribution with rate parameter 1 (i.e., the public scores are zero, and the private scores are exponentially distributed). Let the value of the outside option be $\frac{\log n}{2}$ for every agent. Note that $\mathbb{P}(v_{ij} \geq \frac{\log n}{2}) = \frac{1}{\sqrt{n}}$. Therefore, we expect every agent to have around \sqrt{n} acceptable partners.

Let Alice be given all the workers' preferences, and Bob be given all the firms' preferences. Let π be the communication protocol, at the end of which Alice and Bob output a matching μ_π that is stable with probability at least .9 (that is, the matching protocol is successful with a high constant probability). We claim that it must be the case that the length of the protocol (i.e., the number of bits of communication), is bounded as $|\pi| = \mathbb{E}[|\Pi|] = \Omega(n^{3/2})$, where Π is the realization of the protocol π .

We will focus only on whether a pair of agents find each other mutually acceptable (ignoring other ordinal information), such mutually acceptable pairs being the only pairs that can be matched under any stable matching (this will lead to Claim B.2 below). If worker i and firm j are a mutually acceptable pair, and moreover, this is the only mutually acceptable pair of which each of i and j is a member, then (i, j) *must* be a matched pair under any stable matching. This will yield Claim B.1 below. We will draw on ideas from information complexity theory (see, e.g., Braverman (2015)), together with Claims B.1 and B.2, to establish Claim B.3 and hence our lower bounds. Note that our proof is short and

self-contained, using only basic facts from information theory (Cover and Thomas, 2012).

We define the following boolean random variables for each worker-firm pair (i, j) :

$$A_{ij} := \begin{cases} 1 & \text{if } v_{ij} \geq \frac{\log n}{2}, \\ 0 & \text{otherwise} \end{cases}, \quad \text{and} \quad B_{ij} := \begin{cases} 1 & \text{if } u_{ji} \geq \frac{\log n}{2}, \\ 0 & \text{otherwise.} \end{cases} \quad (8)$$

In other words, $A_{ij} = 1$ means that worker i likes firm j more than the outside option, and $B_{ij} = 1$ means that firm j likes worker i more than the outside option. Note that all A_{ij} , B_{ij} are distributed as Bernoulli(α), where $\alpha = 1/\sqrt{n}$, and are independent of each other. In addition, let M_{ij} be the indicator random variable of whether worker i is matched to firm j under μ_π .

Claim B.1. *For sufficiently large n , we have $\mathbb{P}[M_{ij} = 1 | A_{ij} = B_{ij} = 1] > 10^{-2}$.*

Proof. Assume $A_{ij} = B_{ij} = 1$, so worker i and firm j find one another acceptable. By a standard Chernoff bound, the probability that worker i has more than $2\sqrt{n}$ acceptable partners is at most $\exp(-\frac{\sqrt{n}}{3})$. Since the probability that each of these firms finds worker i to be acceptable is exactly $\frac{1}{\sqrt{n}}$, the probability that another firm out of these other than j finds worker i acceptable is at most $1 - (1 - \frac{1}{\sqrt{n}})^{2\sqrt{n}}$, which converges to $1 - e^{-2}$ for large n . For sufficiently large n , the sum of these two probabilities is no more than $1 - e^{-2.1}$, so the chance that j is the unique firm that both finds i acceptable and also is acceptable to i is at least $e^{-2.1}$. Since the preferences of workers and firms are independent, the chance that both i and j are the unique mutually acceptable partners for one another is at least $e^{-4.2}$ for sufficiently large n . Therefore, for sufficiently large n ,

$$\begin{aligned} & \mathbb{P}[M_{ij} = 1 | A_{ij} = B_{ij} = 1] \\ & \geq \mathbb{P}[\pi \text{ outputs a stable match}] \cdot \mathbb{P}[i \text{ is with } j \text{ in all stable matches} | A_{ij} = B_{ij} = 1] \\ & > (.9) \cdot e^{-4.2} \\ & > 10^{-2}. \end{aligned}$$

□

Claim B.2. $\mathbb{P}[M_{ij} = 1 | A_{ij} = 1, B_{ij} = 0] = 0$, $\mathbb{P}[M_{ij} = 1 | A_{ij} = 0, B_{ij} = 1] = 0$, and $\mathbb{P}[M_{ij} = 1 | A_{ij} = 0, B_{ij} = 0] = 0$.

Proof. This follows from the fact that two agents can be matched with one another in a stable matching only if both find one another acceptable (more preferable than the outside option). □

Note that Claims B.1 and B.2 together imply that π must approximately compute the value of the AND function $A_{ij} \wedge B_{ij}$ in the sense that knowing that $M_{ij} = 1$ implies that $A_{ij} \wedge B_{ij} = 1$, and when $A_{ij} \wedge B_{ij} = 1$, we know that $M_{ij} = 1$ with at least a constant probability. Next, we make the following information-theoretic claim, which quantifies the information complexity of approximately computing the boolean AND function (this line of reasoning is similar to Braverman et al., 2013).

Claim B.3. *We have*

$$I(A_{ij}B_{ij}; \Pi) = \Omega(\alpha) = \Omega(1/\sqrt{n}). \quad (9)$$

Here Π is again the random variable representing the realization of the protocol π , and $I(X; Y)$ is Shannon's mutual information, which, informally, measures the amount of information a random variable X contains about a variable Y (and vice versa). In terms of Shannon's entropy $H(\cdot)$, the mutual information is $I(X; Y) = H(X) - H(X|Y) = H(Y) - H(Y|X)$. In other words, Alice and Bob cannot hope to even approximate the value of $A_{ij} \wedge B_{ij}$ without revealing a substantial amount of information about themselves. Note that fully revealing the values of A_{ij}, B_{ij} corresponds to Shannon's entropy $H(A_{ij}, B_{ij}) = \Theta(\log n/\sqrt{n})$. Let us prove Claim B.3.

Proof. We rely on the following basic facts about protocols and about mutual information.

1. If A_{ij} and B_{ij} are independent, and independent of the players' other inputs, then for each transcript realization π of Π , the variables $(A_{ij}|\Pi = \pi)$ and $(B_{ij}|\Pi = \pi)$ are also independent. The reason is that one player speaking at a time cannot introduce a dependence between these variables (the formal proof is by induction on protocol rounds).
2. Therefore, we have by the Chain Rule (see, e.g., Cover and Thomas (2012) for information theory basics):

$$\begin{aligned} I(A_{ij}B_{ij}; \Pi) &= I(A_{ij}; \Pi) + I(B_{ij}; \Pi|A_{ij}) = I(A_{ij}; \Pi) + I(B_{ij}; \Pi A_{ij}) - I(B_{ij}; A_{ij}) = \\ &= I(A_{ij}; \Pi) + I(B_{ij}; \Pi A_{ij}) = I(A_{ij}; \Pi) + I(B_{ij}; \Pi) + I(B_{ij}; A_{ij}|\Pi) = I(A_{ij}; \Pi) + I(B_{ij}; \Pi). \end{aligned}$$

Therefore, it will be enough to lower bound $I(A_{ij}; \Pi) + I(B_{ij}; \Pi)$.

3. We can write the mutual information expression we are interested in in terms of KL-divergence as follows:

$$I(A_{ij}; \Pi) = \mathbf{E}_{\pi \sim \Pi} \mathbf{D}_{KL}(A_{ij} |_{\Pi=\pi} \| A_{ij}). \quad (10)$$

A similar expression holds for B_{ij} . Again, a proof and further discussion can be found in information theory texts such as Cover and Thomas (2012).

4. It can be shown by direct calculation that for any constant $c < 1$, and $x < 1/2$, it is the case that for $c' < c$

$$\mathbf{D}_{KL}(\text{Bernoulli}(c' \cdot x) \| \text{Bernoulli}(x)) = \Omega_c(x), \quad (11)$$

where the Bernoulli random variable $\text{Bernoulli}(x)$ takes the value 1 w.p. x , and the value 0 w.p. $1 - x$.

By Claims B.1 and B.2 we have that

$$\mathbb{P}[M_{ij} = 0] \geq \mathbb{P}[(A_{ij}, B_{ij}) \neq (1, 1)] = (1 - \alpha^2), \quad (12)$$

and

$$\mathbb{P}[M_{ij} = 0, (A_{ij}, B_{ij}) = (1, 1)] < \alpha^2 \cdot (1 - 10^{-2}). \quad (13)$$

Therefore, for a sufficiently large n (and thus a sufficiently small α),

$$\mathbb{P}[(A_{ij}, B_{ij}) = (1, 1) | M_{ij} = 0] < \frac{\alpha^2(1 - 10^{-2})}{1 - \alpha^2} \leq \alpha^2 \cdot (1 - 9 \cdot 10^{-3}). \quad (14)$$

Let $\Pi_{M_{ij}=0}$ be the distribution of the history of the protocol, conditional on $M_{ij} = 0$, we have by observation 1 above that

$$\mathbf{E}_{\pi \sim \Pi_{M_{ij}=0}} \mathbb{P}[(A_{ij}, B_{ij}) = (1, 1) | \Pi = \pi] = \mathbb{P}[(A_{ij}, B_{ij}) = (1, 1) | M_{ij} = 0],$$

and thus by Markov's inequality

$$\mathbb{P}_{\pi \sim \Pi_{M_{ij}=0}} [\mathbb{P}[(A_{ij}, B_{ij}) = (1, 1) | \Pi = \pi] < \alpha^2 \cdot (1 - 2 \cdot 10^{-3})] \geq 1 - \frac{0.991}{0.998} > 7 \cdot 10^{-3}. \quad (15)$$

Note that $\mathbb{P}[(A_{ij}, B_{ij}) = (1, 1) | \Pi = \pi] < \alpha^2 \cdot (1 - 2 \cdot 10^{-3})$ implies that either $\mathbb{P}[A_{ij} = 1 | \Pi = \pi] < \alpha \cdot (1 - 10^{-3})$ or $\mathbb{P}[B_{ij} = 1 | \Pi = \pi] < \alpha \cdot (1 - 10^{-3})$, and by (11) above, for any $\pi \in \Pi_{M_{ij}=0}$,

$$\mathbf{D}_{KL}(A_{ij} |_{\Pi=\pi} \| A_{ij}) + \mathbf{D}_{KL}(B_{ij} |_{\Pi=\pi} \| B_{ij}) = \Omega(\alpha). \quad (16)$$

By (12) and (15), the probability of such a π is at least $7 \cdot 10^{-3} \cdot (1 - \alpha^2) > 6 \cdot 10^{-3}$ for sufficiently large n . The contribution of such π 's to the expectation of $\mathbf{D}_{KL}(A_{ij}|\Pi=\pi\|A_{ij}) + \mathbf{D}_{KL}(B_{ij}|\Pi=\pi\|B_{ij})$ is therefore at least $\Omega(\alpha)$, since \mathbf{D}_{KL} is always nonnegative, which implies by (10) that

$$I(A_{ij}; \Pi) + I(B_{ij}; \Pi) = \mathbf{E}_{\pi \sim \Pi} [\mathbf{D}_{KL}(A_{ij}|\Pi=\pi\|A_{ij}) + \mathbf{D}_{KL}(B_{ij}|\Pi=\pi\|B_{ij})] = \Omega(\alpha),$$

concluding the proof. \square

Fact B.4. *If X and Y are independent, then $I(X, Y; Z) \geq I(X; Z) + I(Y; Z)$.*

We can now conclude the proof of Theorem 2. Since A_{ij}, B_{ij} are mutually independent for the different values of (i, j) , we get using the above fact, and the fact that entropy is always an upper bound on mutual information,

$$H(\Pi) \geq I(A_{11}B_{11} \dots A_{nn}B_{nn}; \Pi) \geq \sum_{i=1}^n \sum_{j=1}^n I(A_{ij}B_{ij}; \Pi) = n^2 \cdot \Omega(1/\sqrt{n}) = \Omega(n^{3/2}).$$

(We have used Claim B.3 here.) Observing that $|\pi| \geq H(\Pi) = \Omega(n^{3/2})$ concludes the proof of the lower bound on communication cost.

The preference learning lower bound follows, since if there is a protocol of preference learning cost R , each time a preference oracle call is made, Alice (or Bob) can share the learned preference with the other player at cost $O(\log n)$, yielding a communication protocol with cost $C = O(R \log n)$. Therefore $R = \Omega(n^{3/2}/\log n)$.

C A generic algorithm for stable matchings in tiered markets

In this section, we present a generic algorithm for computing stable matchings in tiered markets, based on the ability to compute stable matchings between single tiers of workers and firms. One corollary of this is that the CEDA protocol in Section 3.1 can be generalized to *tiered separable* markets, which are tiered markets in which the preferences between any tier of workers and any tier of firms follow assumptions of the separable market as in Section 2.4. This implies that the $O^*(\sqrt{n})$ average communication cost of computing a stable matching can also be attained for generalizations of separable markets that allow for arbitrarily many

tiers. (Recall that the assumptions on separable markets in Section 2.4 restrict it to having only a constant number of tiers.)

Before presenting the algorithm (Theorem 7), we first define the concept of sub-matching. For any matching μ , any subset $A \subseteq I$ of workers and $B \subseteq J$ of firms, let $\mu(A, B)$ be the *sub-matching* restricted to agents $A \times B$, which is defined as $\{(i, j) \in \mu : i \in A, j \in B\}$. We say that the sub-matching $\mu(A, B)$ is stable if everyone prefers to be matched to their partner over being unmatched and there are no blocking pairs in $A \times B$.

Theorem 7. *A stable matching in a tiered market can be constructed as follows. Initialize $\mu = \emptyset$.*

1. *Construct a stable sub-matching between the top tiers I_1 and J_1 . Add these matches to μ .*
2. *Remove any matched agent in the sub-matching found above, as well as any unmatched agent in I_1 or J_1 who finds someone in J_1 or I_1 unacceptable. (These agents will find all agents in worse tiers unacceptable.) After this, either I_1 or J_1 would have been completely removed.*
3. *If either all of the workers or all of the firms have been removed, then return μ . Otherwise repeat step 1 for the top remaining tiers on both sides.*

Moreover, every stable matching can be constructed in the above way.

Proof. The theorem follows from the following claim, which implies that the set of stable matchings in tiered markets can be decomposed into the Cartesian product of stable sub-matchings for the top tiers and stable sub-matchings for the rest of market.

Claim C.1. *In a tiered market, a matching μ is stable if and only if*

1. *Sub-matching $\mu(I_1, J_1)$ is stable.*
2. *Sub-matching $\mu(I_1 \setminus (I_1^m \cup I_1^u), J_1 \setminus (J_1^m \cup J_1^u))$ is stable, where $I_1^m \subseteq I_1$ denotes the matched workers in $\mu(I_1, J_1)$, and $I_1^u \subseteq I_1$ denotes the unmatched workers who find someone in J_1 unacceptable. The sets of firms J_1^m and J_1^u are similarly defined.*

Given this claim, both directions of Theorem 7 follow from straightforward induction. To show the first direction of this claim, assume that μ is a stable matching for the tiered

market. Note that in any stable matching μ , for a fixed set A of workers and fixed set B of firms, the sub-matching $\mu(A, B)$ must be stable. Therefore, sub-matching $\mu(I_1, J_1)$ must be stable. Apply the Rural Hospital Theorem on this sub-market, we have that the sets I_1^m , I_1^u , J_1^m and J_1^u are fixed in all possible stable matchings μ . Hence, the sets $I_1 \setminus (I_1^m \cup I_1^u)$ and $J_1 \setminus (J_1^m \cup J_1^u)$ are fixed and the sub-matching $\mu(I_1 \setminus (I_1^m \cup I_1^u), J_1 \setminus (J_1^m \cup J_1^u))$ must be stable.

For the second direction of the claim, suppose that the designated sub-matchings are stable, we show that μ must be stable. To do this, we need to show that workers in I_1^u and firms in J_1^u cannot be matched in any stable matching, and that there can be no blocking pairs between workers I_1^m and any firm in J , and no blocking pairs between firms J_1^m and any worker in I .

First, observe that workers I_1^u are unmatched in the stable sub-matching $\mu(I_1, J_1)$, which implies that these workers cannot be matched to anyone in J_1 in a stable matching for the whole market. However, because they find certain firms in the top tier J_1 unacceptable, they must find every firm in worse tiers unacceptable, so cannot be matched to them either. A similar statement can be made for firms in J_1^u .

Now, there can be no blocking pairs between workers in I_1^m and firms in J_1 by the fact that sub-matching $\mu(I_1, J_1)$ is stable. Moreover, there cannot be any blocking pairs between workers in I_1^m and firms in $J \setminus J_1$ because the worker is already matched to someone of a better tier. This implies that there cannot be blocking pairs between workers in I_1^m and any firm. A similar statement can be made for firms in J_1^m . This implies that μ is stable, as desired. \square

Corollary C.2 (Generalization of the Rural Hospital Theorem). *In a tiered market, if we define the partner tier of a given agent in a matching as the tier index of the agent's matched partner (and zero if the agent is unmatched), then for every agent, the partner tier of that agent is the same in every stable matching.*

D Proof of correctness and efficiency of the targeted signaling protocol (Theorem 3)

In this section, we prove Theorem 3, which claims that the targeted signaling protocol succeeds with high probability and bounds its communication and preference learnings costs. The proof is based on studying the properties of the following mathematical object.

Definition D.1. In the targeted-signaling protocol, define the subgraph of signals as the collection of tuples (i, j) for which either

1. The worker i signaled to firm j in the signaling round.
2. The firm j signaled to the worker i and $s_{k(i)} \neq t_{l(j)}$. (We do not count signals from firms who are at an equal position with their target tier, $s_{k(i)} = t_{l(j)}$.)

We break the proof of Theorem 3 into 3 claims.

- **Claim 1:** With probability at least $1 - \frac{18}{n}$, the subgraph of signals contains a stable matching.
- **Claim 2:** Whenever the subgraph of signals contains a stable matching, the matching returned by the targeted signaling protocol is stable (with respect to complete preferences).
- **Claim 3:** The total number of signals is at most $\Theta(n \log^3 n)$.

Claims 1 and 2 imply that the targeted signaling protocol succeeds with probability at least $1 - \frac{18}{n}$. Claim 3 implies the desired bounds on communication and preference learning costs. This is because in the signaling round, sending each signal requires $O(\log(n))$ communication cost and $O(1)$ preference learning cost. In the matching round, the sum of the length of everyone's partial rankings is exactly twice the total number of signals, and producing each ranking of length k requires $O(k \log(n))$ communication cost and $O(k)$ preference learning cost.

D.1 Proof of Claim 1

Definition D.2. Define the tiered DA matching as the matching produced when running the algorithm from Theorem 7 on the tiered random market, with the stable sub-matching between top tiers I_1 and J_1 in step 1 being produced by the following algorithm: if $|I_1| \leq |J_1|$, run the worker-proposing DA algorithm in the sub-market with only tiers I_1 and J_1 ; otherwise, run the firm-proposing DA algorithm in this sub-market.

We prove Claim 1 by proving that with probability $1 - \frac{18}{n}$, the tiered DA matching is contained in the subgraph of signals. The crux is proving the following lemma, which gives a

bound on the rank obtained by a given agent in a uniformly random matching market with certain partners being unavailable. (Intuitively, the unavailable partners represent those that have been matched to better tiers in the tiered DA matching). As in the whole paper, \log here denotes the natural logarithm.

Lemma D.3. *Consider a matching market with m workers, $n \geq m$ available firms and u unavailable firms. The preferences of workers for the $n + u$ firms are uniformly random, and the preference of available firms for workers are uniformly random. The unavailable firms prefer to be unmatched. Let $N \geq 2$ be such that $N \geq n + u$. For any given worker, with probability at least $1 - \frac{9}{N^2}$, we have that in the worker-proposing DA algorithm, the given worker is matched to one of his top r firms, where*

$$r = 24 \frac{n + u}{n} \log^2(N).$$

Proof of Lemma D.3. First, note that without loss of generality, $N \geq 100$ because otherwise, $r > N$.

Label the fixed worker to be worker 1. Label the available firms 1 through n , and the unavailable firms $n + 1$ through $n + u$. Consider the firm-optimal stable match in the sub-market without worker 1 and without the unavailable firms, and call this matching μ_1 . Note that μ_1 does not depend on worker 1's preferences, nor does it depend on the preference of firms for worker 1. Define E_1 as the event that the total rank of workers in matching μ_1 (ignoring the unavailable firms) does not exceed $R = 4e(m - 1) \log(N)$. We lower bound the probability of E_1 using a proposition we prove in Appendix E about the average rank of workers in this setting (Proposition E.1). By plugging in $z = 2 \log(N)$ into Proposition E.1, we have that the probability of event E_1 is at least $1 - \frac{8}{N^2}$.

Let μ_2 be the matching formed by running the worker-proposing DA algorithm from initialization μ_1 . In other words, suppose that we start with everyone else matched according to μ_1 and have worker 1 propose to his top choice as in the DA algorithm. This may cause a previously matched worker to be rejected from a firm, and we will have this worker apply to his next choice, which may result in a chain of rejections leading to someone applying to one of the $n - m + 1$ unmatched available firms. μ_2 is a stable matching (with respect to the entire market). Because the rank of worker 1 in μ_2 is no better than in the worker-optimal stable match, it suffices to upper-bound the rank obtained by worker 1 in μ_2 .

First, let us make a few structural observations on μ_1 and μ_2 under event E_1 . For each firm $j \leq n$, let B_j be the set of workers who weakly prefer firm j to their partner in μ_1 . (We use the letter B because this is the set of workers who want to block with j in μ_1 .) In μ_1 , firm j is matched to the firm's favorite worker in B_j . Furthermore, the sum $\sum_{j=1}^n |B_j| \leq R = 4e(m-1) \log(N)$, as this sum always equals the total rank obtained by workers in μ_1 (ignoring the unavailable firms). Now, consider running the DA algorithm with initialization μ_1 , drawing only as needed the preference of firms for worker 1. When worker 1 applies to an available firm j , the probability that he is accepted is exactly $p_j = \frac{1}{1+|B_j|}$, because this is his chance of being the firm's favorite worker among a set of size $1 + |B_j|$. If the worker is rejected, then the worker apply to his next choice, and the same formula for acceptance probability will apply. If the worker is accepted, then this will trigger a rejection chain that ends with one of the $n - m + 1$ unmatched available firms. Note that this rejection chain can never circle back to firm j and cause worker 1 to be rejected, because that would contradict the assumption that μ_1 is the firm-optimal stable match in the sub-market without worker 1. For the unavailable firms $j \geq n$, define $p_j = 0$. We have that by Jensen's inequality,

$$\sum_{j=1}^{n+u} p_j = \sum_{j=1}^n p_j \geq \frac{n}{4e \log N + 1}. \quad (17)$$

Now, conditional on event E_1 , let P be the probability that worker 1 is not matched to his top r firms in the worker-proposing DA algorithm. Let A be all subsets of the $n+u$ firms of cardinality $\lfloor r \rfloor$. We have,

$$P \leq \frac{1}{|A|} \sum_{S \in A} \prod_{j \in S} (1 - p_j) \quad (18)$$

$$\leq \left(1 - \frac{\sum_{j=1}^n p_j}{n+u} \right)^r \quad (19)$$

$$\leq \exp \left(- \frac{nr}{(n+u)(4e \log N + 1)} \right) \quad (20)$$

$$< \frac{1}{N^2} \quad (21)$$

Inequality (18) follows from the independence between the preference of worker 1, the event E_1 , and each firm j 's preference for worker 1. Inequality (19) follows from the fact that the sum of product in the inner part of equation 18 increases if we replace any two different p_i, p_j with their average $\frac{p_i+p_j}{2}$, so the maximum is attained when all of them are equal. Inequality

(20) follows from inequality 17 and the bound $(1-x) \leq \exp(-x)$, and inequality (21) follows from the fact that when $N \geq 100$, we have $\frac{24 \log^2 N}{4e \log N + 1} > 2 \log N$.

Since the probability that E_1 does not occur is at most $\frac{8}{N^2}$, we have that the total probability that worker 1 does not get one of his top r choices is at most $\frac{8}{N^2} + \frac{1}{N^2} = \frac{9}{N^2}$, which is what we needed. \square

Claim 1 follows from Lemma D.3 and taking an union bound over the $m + n \leq 2n$ agents, observing that in the tiered DA matching, every agent is in a scenario described by Lemma D.3. So with probability at least $1 - \frac{18}{N}$, the tiered DA matching is contained in the subgraph of signals.

D.2 Proof of Claim 2

We break the proof of Claim 2 into two parts.

1. **Claim 2a)** Whenever the subgraph of signals contains a stable matching, running worker-proposing DA with preferences restricted to this subgraph returns a matching μ that is stable with respect to complete preferences.
2. **Claim 2b)** Whenever the above happens, the targeted signaling protocol returns the same matching μ .

We first show Claim 2b). Let μ denote the matching that arises from running worker-proposing DA with preferences restricted to the subgraph of signals. If μ is stable (with respect to complete preferences), then it must match every worker i who has a target tier: *i.e.* there exists firm tier l such that $t_l \geq s_{k(i)}$. Now, observe that the output of the targeted signaling protocol is simply running the worker-proposing DA with respect to the subgraph of signals plus certain additional edges. Precisely speaking, these edges are the ones ruled out in Definition D.1, which are tuples (i, j) where j signaled to i and they $s_{k(i)} = t_{l(j)}$. Since the DA algorithm without these edges ended up matching every one of these agents i incident to one of these edges, and the favorite partner of agent i are already included in the subgraph of signals, running DA with these edges would not change the result. Therefore, the targeted signaling protocol returns μ .

Claim 2a) follows from the following structural result (Lemma D.4) on the subgraph of signals. First, let us give a few definitions. For any subgraph (defined as a collection of worker-firm tuples (i, j)), we say that a matching μ is *stable with respect to the subgraph* if

every matched agent in μ prefers to be matched than unmatched, and there are no blocking pairs to μ within the subgraph. Define the *complete graph* as the Cartesian product $I \times J$. (The original definition of stability is equivalent to stability with respect to the complete graph.) Define a matching to be *full* if it has cardinality $\min(n_I, n_J)$, which corresponds to matching all agents of at least one of the sides.

Lemma D.4. *Any matching μ that is full and stable with respect to the subgraph of signals is stable with respect to the complete graph.*

To see why Claim 2a) follows from Lemma D.4, note that when the subgraph of signals contains a stable matching, then running worker-proposing DA with preferences restricted to the subgraph returns a matching that is full and that is stable with respect to the subgraph. Lemma D.4 implies that this matching is stable with respect to the complete graph.

Proof of Lemma D.4. Let the subgraph of signals be G . let μ be a full matching stable with respect to G , we show that μ is stable with respect to the complete graph. It suffices to show that μ does not contain any blocking pairs.

We show that no tuple (i, j) in the complete graph can be a blocking pair. Let worker i be in worker-tier k and firm be j in firm-tier l . Firstly, any $(i, j) \in G$ cannot be a blocking pair by the definition of μ being stable in G .

Suppose first that $s_k \neq t_l$. Without loss of generality, let $s_k < t_l$. In this case, worker i must be matched in μ , say to firm j' . Let \tilde{l} be the target tier of worker i and let l' be the tier of firm j' . Note that $l' \leq \tilde{l}$ because G can only contain edges between i and weakly better tiers than \tilde{l} . Moreover, we have $\tilde{l} \leq l$, since $\tilde{l} = \min\{l : t_l \geq s_k\}$ by definition. Combining the two inequalities, we have $l' \leq \tilde{l} \leq l$. Suppose that $l' < l$, then i would prefer j' to j , so (i, j) cannot be a blocking pair. Suppose that $l' = l$, then both must equal \tilde{l} , which is the target tier of workers I_k . This means that the target tier of firms J_l must be a strictly worse tier of workers than I_k . (Otherwise, we would need $s_k = t_l$.) So the only reason that $(i, j') \in \mu$ is in the subgraph of signals G is that i signals to j' . But i does not send a signal to j , and both j and j' are in the same tier, so i must prefer j' to j , so (i, j) cannot be a blocking pair.

The only remaining case is $s_k = t_l$. In this case, the subgraph of signals (Definition D.1) only includes signals from I_k to J_l . Thus, i must be matched in μ , say to j' , and we have that j' is either in a better tier than j or is signaled to by i . In either cases, i must prefer j' to j , so (i, j) cannot be a blocking pair.

Since there are no blocking pairs, μ must be stable with respect to the complete graph. \square

D.3 Proof of Claim 3

We complete the proof of Theorem 3 by proving Claim 3.

Lemma D.5. *The total number of signals sent during the signaling round of the targeted signaling protocol is at most $60n \log^3 n$.*

Proof of Lemma D.5. The result is trivially true if $n \leq e^4 < 60$, since the average number of signals is at most n . Assume now that $\log n \geq 4$.

Define a *grouping* of tiers as the set of all tiers of agents that share the same target tier. For example, if the shared target tier is firm-tier J_l , then the grouping is all worker-tiers I_k such that $t_{l-1} < s_k \leq t_l$.

Consider an arbitrary grouping. Without loss of generality, let the shared target tier be J_l as above. Let $K = \{k : t_{l-1} < s_k \leq t_l\}$ as above. Let the minimum and maximum element of K be k_0 and k_1 respectively. For each $k \in K$, define r_k as the target number of workers in tier I_k (see Definition 4.4).

Let the total number of signals sent by this grouping be $\sigma = \sum_{k=k_0}^{k_1} m_k r_k$. Note first that $r_{k_0} = 24 \log^2 n$ since the competitiveness of workers (see Definition 4.3) in I_{k_0} must be 1. For $K \setminus \{k_0\}$, we have

$$\begin{aligned} \sum_{k=k_0+1}^{k_1} m_k r_k &= 24n_l \log^2 n \sum_{k=k_0+1}^{k_1} \frac{m_k}{t_l - s_{k-1}} \\ &\leq 24n_l \log^2 n \sum_{k=k_0+1}^{k_1} \left(\frac{1}{t_l - s_{k-1}} + \cdots + \frac{1}{t_l - s_k + 1} \right) \\ &\leq 24n_l \log^2 n \left(\frac{1}{t_l - s_{k_0}} + \frac{1}{t_l - s_{k_0} - 1} + \cdots + \frac{1}{2} + \frac{1}{1} \right) \\ &\leq 24n_l \log^2 n \log(t_l - s_{k_0} + 1) \\ &\leq 24n_l \log^3 n \end{aligned}$$

This shows that

$$\sigma \leq 24m_{k_0} \log^2 n + 24n_l \log^3 n \leq (6m_{k_0} + 24n_l) \log^3 n.$$

From this the desired result follows because the sum of all possible m_{k_0} is at most $2n$, and the same holds for the sum of all possible n_l . \square

E Average rank in unbalanced matching market

The proof of Lemma D.3 requires a concentration bound for the average rank obtained by workers in any stable matching in a uniformly random matching market with strictly more firms than workers. When there are $n - 1$ workers and n firms, Ashlagi et al. (2017) implies that this is asymptotically $\log(n)$ as $n \rightarrow \infty$. Compared to their result, the following proposition gives up a constant factor of $2e$ in the asymptotics, but obtains a stronger probabilistic guarantee that holds for any n . The proof builds on the techniques from Pittel (1989) and Pittel (1992), which are based on an Integral formula due to Knuth.

Proposition E.1. *Consider a uniformly random matching market with $n - 1$ workers and n firms. For any $z \geq 2$, we have that with probability at least $1 - 8 \exp(-z)$, the average rank obtained by workers in any stable matching is no more than*

$$\bar{r} = e(2 \log n + z).$$

Proof. Define $m = n - 1$. Let P_0 be the probability that a uniformly random matching market with m workers and n firms has a stable matching in which the average rank of workers is greater than \bar{r} . We upper-bound P_0 by $8 \exp(-z)$. Observe that this is trivially true if $m \leq 7$ because $\bar{r} \geq 3e > 8$, so we assume from now on that $m \geq 8$.

As in Pittel (1989), define the standard matching μ_0 as the matching in which for each $1 \leq i \leq m$, worker i is matched to firm i . The unmatched firms are denoted by indices $j > m$. For each tuple (i, j) , $1 \leq i \leq m$, $1 \leq j \leq n$, let X_{ij} and Y_{ij} be i.i.d. draws from $\text{Uniform}[0, 1]$. These values induce the preferences of workers and firms as follows: the smaller the value of X_{ij} , the more worker i prefers firm j . Similarly, the smaller the value of Y_{ij} , the more firm j prefers worker i . For simplicity, define $x_i = X_{ii}$ and $y_i = Y_{ii}$ for $1 \leq i \leq m$. We call the matrices X and Y the cardinal preferences of workers and firms, and the vectors x and y the matching values of workers and firms.

Adapting Equation (2.2) of Pittel (1989), we have that given the matching values x and y , the probability the standard matching is stable and the total rank of workers equals R is exactly

$$[\xi^{R-m}] \left\{ \prod_{i=1}^m (1 - x_i) \prod_{1 \leq i \neq j \leq m} (1 - x_i + \xi x_i (1 - y_j)) \right\}, \quad (22)$$

where $[\xi^a]\{f(\xi)\}$ denotes the coefficient of ξ^a in the expansion of polynomial $f(\xi)$. This formula is analogous to Equation (2.2) of Pittel (1989), and we provide a brief explanation below. The rank obtained by each worker is exactly one plus the number of firms the worker wants to block with, so the total rank of workers is R if and only if the total number firms workers want to block with is $R - m$, counting with multiplicity. The expression in the braces computes the probability the standard matching μ_0 is stable while keeping track of who wants to block with whom using dummy variable ξ . The expression is a product of various terms, and is based on the i.i.d. assumptions of entries of X and Y . In the first product, $(1 - x_i)$ is the probability that worker i does not want to block with the unmatched firm, as $P(X_{in} > X_{ii}) = 1 - x_i$. In the second product, we have the linear combination of two terms: $1 - x_i$ is the probability that worker i does not want to block with firm j , and $x_i(1 - y_j)$ is the probability that worker i wants to block with j but j does not reciprocate. Expanding the product as a polynomial in ξ and examining the coefficient of ξ^{R-m} obtains exactly the probability that the matching is stable and the total rank of workers is R .

Define $A = \{(x, y) : 0 \leq x_i, y_i \leq 1, 1 \leq i \leq m\}$. We have

$$P_0 \leq n! \int_A \sum_{R=\lfloor m\bar{r} \rfloor + 1}^{\infty} [\xi^{R-m}] \left\{ \prod_{i=1}^m (1 - x_i) \prod_{1 \leq i \neq j \leq m} (1 - x_i + \xi x_i(1 - y_j)) \right\} dx dy \quad (23)$$

$$\leq n! \int_A \inf_{\xi \geq 1} \left\{ \xi^{m(1-\bar{r})} \exp(-m \sum_{i=1}^m x_i + \xi \sum_{1 \leq i \neq j \leq m} x_i(1 - y_j)) \right\} dx dy \quad (24)$$

Inequality (23) follows from integrating equation (22) over the uniform distribution of matching values, then using a union bound over all $n!$ matchings based on symmetry. Inequality 24 comes from the Chernoff method of bounding the tail of power series⁴⁶, and using the fact that $\lfloor m\bar{r} \rfloor + 1 \geq m\bar{r}$.

Define $\mu = \frac{\bar{r}}{e} = 2 \log(n) + z$. Partition the region of integration A into $A_1 = \{(x, y) \in A : \sum_{i=1}^m x_i \geq \mu\}$, and $A_2 = A \setminus A_1$. Define the above integral in regions A_1 and A_2 to be P_1 and P_2 respectively. It suffices to bound P_1 and P_2 . For convenience, define $s = \sum_{i=1}^m x_i$.

To bound P_1 , we set $\xi = 1$. Let $s_i = s - x_i$, $\Psi(x) = \int_0^1 \exp(-xy) dy = \frac{1 - \exp(-x)}{x}$, and $A_1^x = \{x : 1 \leq x_i \leq 1, \sum_{i=1}^m x_i \leq \mu\}$. For clarity, we write a series of inequalities and explain

⁴⁶For any power series $f(\xi)$ with positive coefficients, $\sum_a^\infty [\xi^a]\{f(\xi)\} \leq \xi^{-a} \inf_{\xi \geq 1} \{f(\xi)\}$.

them one by one afterward.

$$P_1 \leq n! \int_{A_1} \exp(-s - \sum_{i=1}^m s_i y_i) dx dy \quad (25)$$

$$= n! \int_{A_1^x} \exp(-s) \prod_{i=1}^m \Psi(s_i) dx \quad (26)$$

$$\leq e^2 n! \int_{A_1^x} \exp(-s) \frac{1}{s^m} dx \quad (27)$$

$$\leq e^2 n! \int_{\mu}^m \exp(-s) \frac{1}{s^m} \frac{s^{m-1}}{(m-1)!} ds \quad (28)$$

$$\leq e^2 n(n-1) \int_{\mu}^m \exp(-s) dx \quad (29)$$

$$\leq e^2 \exp(-z) \quad (30)$$

Inequality (25) follows from plugging in $\xi = 1$ into inequality (24). In equation (26), we integrate with respect to each y_i . In inequality (27), we make use of the fact that for each $a > 0$,

$$(\log \Psi(a))' = \frac{1}{\exp(a) - 1} - \frac{1}{a} \geq -\frac{2}{a+2}.$$

So

$$\begin{aligned} \log \left(\prod_{i=1}^m \Psi(s_i) \right) &\leq 2 + \sum_{i=1}^m \left(\log(\Psi(s_i)) - \frac{2x_i}{s+1} \right) \\ &\leq 2 + \sum_{i=1}^m \left(\log(\Psi(s_i)) - \frac{2x_i}{s_i+2} \right) \\ &\leq 2 + \sum_{i=1}^m \log(\Psi(s_i + x_i)) \\ &= 2 + m \log(\Psi(s)) \\ &\leq 2 - m \log(s) \end{aligned}$$

In inequality (28), we use a standard change of variable from x to s (see equation (3.2) and inequality (3.3) of Pittel (1989)). In inequality (29), we use the fact that $\frac{1}{s}$ is decreasing in s and that $\mu \geq 1$. In inequality (30), we integrate out s and make use of the definition of μ and $z \geq 1$.

To bound P_2 , we set $\xi = \frac{e\mu}{s}$. As before, we state a series of inequalities and explain them

afterward.

$$P_2 \leq n! \int_{A_2} \left(\frac{e\mu}{s}\right)^{m(1-e\mu)} \exp(-ms + e(m-1)\mu) dx dy \quad (31)$$

$$\leq n^2 \int_0^\mu \left(\frac{e\mu}{s}\right)^{m(1-e\mu)} \exp(-ms + e(m-1)\mu) s^{m-1} dx dy \quad (32)$$

$$\leq n^2 \int_0^\mu \exp(-(m+e)\mu + m)\mu^{m-1} ds \quad (33)$$

$$= [n^2 \exp(-\mu)] [\exp(-(m+e-1)\mu)\mu^m] \exp(m) \quad (34)$$

$$\leq \frac{1}{e} \exp(-z) \quad (35)$$

Equation (31) substitutes in $\xi = \frac{e\mu}{s}$ to inequality (24) and uses the fact that $\sum_{1 \leq i \neq j \leq m} x_i(1-y_j) \leq s(m-1)$. Inequality (32) again integrates out each y_i and uses the change of variable from x to s as in inequality (3.3) of Pittel (1989). Inequality (33) uses the fact that the function $f(s) = \exp(-ms)s^{em\mu-1}$ is increasing in $[0, \mu]$. Equation (34) simplifies the formula and arrange into groups, denoted by square brackets. Inequality 35 comes from bounding each group. We bound the first group by $\exp(-z)$ using the formula for μ . We bound the second group using the observation that the function $f(\mu) = \exp(-(m+e-1)\mu)\mu^m$ is maximized when $\mu = \frac{m}{m+e-1}$, so $f(\mu) \exp(m) \leq \left(\frac{m}{m+e-1}\right)^m = \left(1 - \frac{e-1}{m+e-1}\right)^m \leq \exp\left(-\frac{(e-1)m}{m+e-1}\right) < \exp(-1)$, since $m \geq 8$.

Combining inequalities (30) and (35), we have

$$P_0 = P_1 + P_2 \leq \left(e^2 + \frac{1}{e}\right) \exp(-z) < 8 \exp(-z),$$

which completes the proof. \square

Remark E.2. *The above proof can be modified to show the following statement. Consider a uniformly random matching market with m men and $n = m + d$ women, where $1 \leq d \leq (e-1)m$. For any $z \geq 2$, with probability at least $1 - 8 \exp(-z)$, the average rank of men in any stable matching is no more than*

$$\bar{r} = e \left(\log \left(\frac{n}{d} \right) + \frac{\log m + z}{d} + 1 \right).$$

F Proofs of incentive properties (Theorems 4 and 5)

In a matching protocol, we say that the agent *complies* with the protocol if she truthfully participates according to what is prescribed, and that the agent *deviates* from the protocol otherwise.

Before proving that both protocols in this paper are as incentive compatible as DA with high probability (Theorem 4), we establish two lemmas. This allows us to prove the desired properties for both protocols simultaneously under one framework.

Define the subgraph of signals in the targeted-signaling protocol as in Definition D.1 in Appendix C. For CEDA, define the subgraph of signals as follows.

Definition F.1. *In the CEDA protocol, for a given realization of preferences, define the subgraph of signals as the set of tuples (i, j) for which worker i applies to j at some point during the protocol.*

Let G be the subgraph of signals in either protocols, assuming everyone complies. In either protocols, define the worker-optimal stable match *restricted to subgraph G* as the result of running the worker-proposing DA algorithm using only preference information within pairs of agents in G . Similarly define the firm-optimal stable match restricted to G . Furthermore, we say that a matching is stable with respect to G if it is individually rational⁴⁷ and there are no blocking pairs $(i, j) \in G$, with respect to true preferences. A matching is stable with respect to complete preferences if it is stable with respect to the complete graph $I \times J$.

Lemma F.2. *In either CEDA (Protocol 2) or targeted signaling (Protocol 3), let the subgraph of signals be G , the outputted matching μ is not blocked by any edge $(i, j) \in G$.*

1. *The worker-optimal stable match restricted to G (defined assuming everyone complies) is stable with respect to complete preferences.*
2. *The firm-optimal stable match restricted to G is also a stable with respect to complete preferences.*

Proof of Lemma F.2. In CEDA, the outputted matching μ is simply the result of running the worker-proposing DA on G , so μ is not blocked by any edge G . In the targeted signaling protocol, the outputted matching is the result of running the worker-proposing DA on a graph that contains G ,⁴⁸ so μ must not be blocked by any edge of G .

⁴⁷A matching is individually rational if agents are only matched to partners they find acceptable.

⁴⁸Specifically, the graph is the collection of tuples (i, j) in which at least one signaled to the other during the signaling round. By Definition D.1, this is $G \cup \{(i, j) : j \text{ signaled to } i, s_{k(i)} = t_{l(j)}\}$.

In CEDA, with high probability, the outputted matching is stable with respect to complete preferences (Theorem 1). When this happens, the worker-optimal stable match restricted to G , which is μ , is stable. The firm-optimal stable match restricted to G is also μ . (To see this, note that by definition of G , every firm in μ gets their favorite partner in G .)

In the targeted signaling protocol, with high probability, the subgraph of signals contains a stable matching by Claim 1 of Appendix C. When this happens, the result of the worker-optimal and firm-optimal DA must be full⁴⁹ and stable with respect to the subgraph of signals. Therefore, by Lemma D.4, both matchings are stable with respect to complete preferences. \square

We are now ready to prove Theorem 4 and 5.

Proof of Theorem 4. We prove that both protocols are as incentive compatible as DA with high probability. In either protocol, let G be the subgraph of signals when everyone complies.

Suppose on the contrary that a certain agent can unilaterally deviate from the protocol and cause the resultant matching to be μ' , which gives the deviating agent someone than the agent's best stable partner under complete, true preferences. By Lemma F.2, we have that with high probability, the agent obtains a better partner in μ' than what the agent gets in either the worker-optimal or firm-optimal stable match restricted to G (defined above with respect to true preferences). Note that regardless of the agent's deviation, μ' must be individually rational for all agents. We apply the following Blocking Lemma due to Gale and Sotomayor (1985) (The version below is from Roth and Sotomayor (1990), Lemma 3.5).

Lemma F.3 (Blocking Lemma, Gale and Sotomayor (1985)). *Let μ_W be the worker optimal stable match. Let μ be any individually rational matching with respect to strict preferences R and let I' be all workers who prefer μ to μ_W . If I' is non-empty, there is a worker-firm pair (i, j) which blocks μ such that $i \notin I'$ and $\mu(j) \in I'$.*

The strict preferences R we use for the above lemma are the true preferences restricted to G , treating everything else as unacceptable. In other words, worker i 's preference is induced by her true preferences over $\{j : (i, j) \in G\}$, and if $(i, j) \notin G$, then she treats firm j as unacceptable. Similarly defined preferences of firms. The matching we use is the μ' above.

An implication of the lemma is that there exists $(i, j) \in G$ such that

⁴⁹A matching is full if it matches $\min(n_I, n_J)$ pairs of agents

1. neither i or j is the same as the deviating agent, so are assumed to be complying to the protocol;
2. (i, j) blocks μ' (under the true preferences of i and j).

Let G' be the subgraph of signals under the deviating action by the agent. (Regardless of the deviation, this is well defined at the end of both protocols). Since (i, j) blocks μ' , $(i, j) \notin G'$ by Lemma F.2. However $(i, j) \in G$ by construction. We show that this discrepancy can happen with vanishing probability in either protocol, thus proving the desired result. The common idea is that the deviating agent has limited control over the edges of subgraph of signals involving complying agents only.

In the targeted signaling protocol, $(i, j) \in G$ implies that $(i, j) \in G'$ because both i and j send the same set of signals in the signaling round regardless of what the deviating agent does. This proves that targeted signaling is as incentive compatible as DA with high probability.

In the CEDA protocol, $(i, j) \in G$ implies that $(i, j) \in G'$ with high probability. This is because by construction, both i and j prefer each other to their partners in μ' . Suppose on the contrary that $(i, j) \notin G'$, then i skipped firm j in her application decision, which implies that the firm must not have sent her a preference signal and that she must not have publicly qualified for the firm, $a_{ji} < z_j$, where z_j is the qualification requirement of firm j at the end of the protocol. Let x_j be the matched utility of firm j at the end of the protocol. The fact that j prefers i to its matched partner in μ' implies

$$x_j < u_{ji} = a_{ji} + \epsilon_{ji}.$$

Which means that either $\epsilon_{ji} > q_j$ or $z_j > a_{ji} > x_j - q_j$. In other words, either it's the case that $\epsilon_{ji} > q_j$ but firm j does not send a preference signal to i , or it's the case that $z_j > \max\{\underline{a}_j, x_j - q_j\}$ for some firm j at the end CEDA. But since firm j (by construction) complies to the protocol, the first case happens with vanishing probability by Lemma A.1. The second also happens with vanishing probability by the proof Lemma A.2. The key is to note that for a compliant firm, the invariant proved in Lemma A.2 that $z_j \leq \max\{\underline{a}_j, x_j - q_j\}$ holds with high probability by a statistical argument, based on the randomness of private scores, and is independent of the application decisions of agents, and also independent of whatever strategic actions that influence applications. This proves that CEDA is as incentive compatible as DA with high probability. \square

Proof of Theorem 5. By Theorem 4, it suffices to show that under general imbalance and when the number of agents of every tier is large, the probability that any given agent has multiple stable partners is vanishing. By the decomposition result in Theorem 7 in Appendix C, the desired result follows from the following Lemma, which is based on techniques from Ashlagi et al. (2017). \square

Lemma F.4. *There exists a non-increasing function $\delta : \mathbb{N} \rightarrow \mathbb{R}$ such that $\delta(y) \rightarrow 0$ as $y \rightarrow \infty$ with the following property. In a matching market with a single tier of m workers and a single tier of $n \geq m + 1$ firms, the probability that any given agent, conditional on being matched, has multiple stable partners is upper-bounded by $\delta(n)$.*

Proof of Lemma F.4. Consider two cases, suppose that $m \geq \frac{n}{2}$, then by Theorem 3 ii) of Ashlagi et al. (2017), there exists m_0 such that for all $m > m_0$, with probability at least $1 - \exp(-\log^{0.4} n)$, the number of workers with multiple stable partners is no more than $\frac{m}{\log^{0.5} m}$, and the same statement holds for firms. This implies that for $n \geq 2m_0$, the probability that any given worker has multiple stable partners is at most

$$\exp(-\log^{0.4} n) + \log^{-0.5} \left(\frac{n}{2} \right).$$

Similarly, the probability that a firm, conditional on being matched, has multiple stable partners is also upper-bounded by this. Define function $\delta_1 : \mathbb{N} \rightarrow \mathbb{R}$ has $\delta_1(n) = 1$ for $n < 2m_0$ and as the above quantity when $n \geq 2m_0$. Then function δ_1 satisfies the desired result in the region when $m \geq \frac{n}{2}$.

Suppose now that $m < \frac{n}{2}$, then we show that the probability that a given matched agent has multiple stable partners is still small. Consider the outcome of the worker-proposing DA algorithm. Consider any firm that is matched in this worker-optimal stable match (WOSM). Let $u = n - m \geq \frac{n}{2}$. As in Ashlagi et al. (2017), suppose that the firm has multiple stable partners, then the chain of proposals triggered by the firm rejecting its current partner must come back to this firm before going to the u unmatched firms. The chance that this happens is at most $\frac{1}{u}$. Moreover, if the firm has more than two stable partners, then when the firm rejects the second stable partner, the chain of proposals triggered also needs to come back to the firm rather than go to the u unmatched firms. So the number of stable partners of this firm is stochastically dominated by $\text{Geometric}(\frac{2}{n})$. For $n \geq 11$, the expectation of this is less than $\frac{3}{n}$. Thus, the expected total number of worker-firm pairs that can be in a stable match and that is not already in the WOSM is at most $\frac{3m}{n}$ for any $n \geq 11$. By symmetry, for

any agent, conditional on being matched, the chance the agent has multiple stable partners is no more than $\frac{3}{n}$ when $n \geq 11$. Define $\delta_2 : \mathbb{N} \rightarrow \mathbb{R}$ to be $\delta(n) = 1$ for $n < 11$ and $\delta(n) = \frac{3}{n}$, we have that δ_2 satisfies the desired result in the region when $m < \frac{n}{2}$.

Finally, we note that $\delta = \max(\delta_1, \delta_2)$ satisfies the desired property for any $n \geq m + 1$. \square

G Impossibility of Two-Round Protocol for General Separable Markets

In this section, we show a simple example of a separable market for which no two-round protocol that uses $o(n)$ bits of communication per agent computes a stable matching with high probability.

For clarity of exposition, the example is described with certain agents preferring any partner of a specific type over any partner of another type. This can be approximated with high probability with private scores distributed as $\text{Exp}(1)$ by having a public score difference of $3 \log n$ for the type that one prefers.

Example 1. *Consider an academic job market with two types of departments and two types of applicants. The types are teaching-focused departments and applicants, and research-focused departments and applicants. An agent prefers to be matched with a partner of his or her own focus, and otherwise preferences are drawn uniformly at random. Research-focused departments and teaching focused candidates are in short supply: there are n research-focused departments, $n + 2$ teaching-focused departments, $n + 2$ research-focused applicants and n teaching-focused applicants.*

In this example, in any stable matching, there are two research-focused applicants and two teaching-focused departments who are matched with each other. However, a priori, the chance that any two agents of different focus are matched in a stable matching is $\frac{2}{n}$. So in a two-round protocol, it's never worthwhile for agents to signal across their own focus. The communication-efficient method is to first let research-focused departments and teaching-focused applicants pick their partners in a two-round protocol, and then run an additional aftermarket to match the remaining agents.

H Optimal communication cost in tiered random markets

The targeted signaling protocol in Section 4.4 uses $\Theta(\log^4 n)$ bits of communication per agent. In this section, we show that the best possible is $\Theta(\log^2 n)$.

Consider the protocol based on the generic algorithm for tiered markets in Theorem 7, in which in Step 1, if $|I_1| \leq |J_1|$, we simulate the worker-proposing DA algorithm, and if $|I_1| > |J_1|$ then we simulate the firm-proposing. Because we only allow private messages, each time an agent propose to a partner, the agent does not know whether or not the partner is taken by agents from better tiers. This results in wasted proposals. Nevertheless, one can show as in Knuth (1976) that the number of wasted proposals is not too high. In fact, one can show that with high probability, this protocol terminates with a stable matching using $\Theta(n \log n)$ proposals. This shows that there exists a stable matching protocol that only uses private messages and that succeed with high probability, using communication cost $\Theta(\log^2 n)$ per agent and preference learning cost of $\Theta(\log n)$ per agent.

The following shows that this bound on average communication cost is the best possible.

Theorem 8 (Lower bound on communication cost with private messages). *Consider a uniformly random market with $n - 1$ workers and n firms. In such a market, any matching protocol that is stable with high probability and only uses private messages must incur a communication cost of $\Omega(\log^2 n)$ bits per agent.*

The following Theorem says that in certain tiered-random markets, esome agents must experience up to \sqrt{n} bits of communication, even though the average is only poly-logarithmic in n .

Theorem 9 (Lower bound on agent specific communication). *Consider a market consisting of $n-1$ workers, each worker in a tier of her own, and n firms all in a single tier. Consider the worker in the \sqrt{n} -th worst tier. In any matching protocol that is stable with high probability and only uses private messages, this worker must incur a communication cost of at least $\Omega(\sqrt{n})$ bits.*

Proof of Theorem 8. The market in question is the same as that in Ashlagi et al. (2017). It suffices to prove the lower bound for any one worker, since all workers are ex ante the same. Some features of this market that we know from Ashlagi et al. (2017) are:

- In all stable matchings, the average rank of workers for their matched firms is very close to $\log n$, and a vanishing fraction of workers (firms) have multiple stable partners. (The average rank of workers for their matched firms is at least $0.99n/\log n$.)
- Fix a worker i . Whp (with high probability), the worker has a unique stable partner. Conditioned on the stable partner being unique, the distribution of worker i 's rank of her stable partner is asymptotically close to $\text{Geometric}(1/\log n)$. In particular, the conditional probability that the unique stable partner is one of worker i 's top $\log n$ most preferred firms is $p \in (1 - 1/e - 0.01, 1 - 1/e + 0.01)$ for large enough n .
- Run the worker proposing deferred acceptance algorithm with worker i excluded. Whp, all but $n^{0.99}$ firms receive between $0.9 \log n$ and $1.1 \log n$ proposals from workers.

Our proof approach is as follows: we consider some worker i and an oracle who knows the preferences of all other agents, and seek to find, whp, a stable partner of i . A slight complication here is that workers may have multiple stable partners in these markets. We work around this by defining a communication problem P_1 as follows: The correct answer is “Yes” if the unique stable partner of i occurs in her top $\log n$ most preferred firms, and the correct answer is “No” if the unique stable partner of i does not occur in her top $\log n$ most preferred firms. In the case that i does not have a unique stable partner, either a “Yes” or a “No” is considered correct. Note that given a candidate stable partner of i , one can output “Yes” if the stable partner is among i 's top $\log n$ most preferred firms, and a “No” if not. If the candidate stable partner is truly a stable partner, the output is a correct answer. Hence, the problem of producing an output that is correct whp for problem P_1 , is no harder than the problem of finding an agent j who, whp, is a stable partner of i .

Call a firm j “accessible” if it satisfies the following: Fix preferences of all agents except worker i . Suppose worker i moves agent j to the top of her preference list, keeping the rest of her preferences unchanged. Then worker i will be matched to agent j under the worker optimal stable matching. Denote by J_a the set of firms accessible to worker i . Note that whether $j \in J_a$ or not *does not depend on the preferences of worker i* . This follows immediately from the fact that the worker optimal stable matching can be computed using the worker proposing deferred acceptance algorithm, and this algorithm makes no use of worker i 's preferences unless she is rejected by firm j , in which case we already know that $j \notin J_a$. Finally, note that when i has a unique stable partner, this is his most preferred firm in J_a .

We now control the size of set J_a , showing that its size is close to $n/\log n$. We make use of an analysis resembling that in Ashlagi et al. (2017) (with revelation of preferences as needed) and using the third bullet stated above. Consider the worker optimal stable matching with worker i excluded. Let J' be the set of firms that have each received between $0.9 \log n$ and $1.1 \log n$ proposals. Using the third bullet, we have

$$|J - J'| \leq n^{0.99}$$

For each of these firms, independently, the probability that they prefer i over their currently matched worker is at least $1/(1 + 1.1 \log n) \geq 0.8/\log n$ and at most $1/(1 + 0.9 \log n) \leq 1.2/\log n$. Let J'' be the set of firms in J' who prefer i over their currently matched worker. It follows using a standard concentration bound that whp, we have

$$0.7n/\log n \leq |J''| \leq 1.3n/\log n.$$

The probability that a rejection chain starting at a firm in J'' , cf. Ashlagi et al. (2017), will return to the firm with a proposal that the firm prefers over worker i , before it terminates by going to the unmatched firm is at most $1/(2 + 0.9 \log n) \leq 1.2/\log n$. Call the set of firms for which the rejection chain returns \hat{J} . These firms may or may not be in J'' . With high probability, using Markov's inequality, we have $|\hat{J}| \leq f_n |J''| 1.2/\log n$, for any $f_n = \omega(1)$. Using $f_n = \sqrt{\log n}$ we obtain a bound of

$$|\hat{J}| \leq \sqrt{n} 1.3n/\log n \cdot 1.2/\log n \leq 2n/(\log n)^{3/2}.$$

All firms in $J'' \setminus \hat{J}$ (here the rejection chain terminates without returning to the firm) are for sure a part of J_a . Thus, we have, whp,

$$|J_a| \geq |J'' \setminus \hat{J}| = |J''| - |\hat{J}| \geq 0.7n/\log n - 2n/(\log n)^{3/2} \geq 0.5n/\log n.$$

On the other hand, we have $J_a \subseteq J'' \cup (J - J')$, leading to, whp,

$$|J_a| \leq |J''| + |J - J'| \leq 1.3n/\log n + n^{0.99} \leq 1.5n/\log n.$$

Let the set of the worker's most preferred $\log n$ firms be J_p . Now again consider the communication problem P_1 and take any protocol that solves it. In cases where worker i has a unique stable partner and the protocol finds a correct answer, the answer is exactly $\mathbb{I}(J_a \cap J_p \neq \emptyset)$. Recalling that whp, worker i has a unique stable partner, and since the

output of the protocol matches $\mathbb{I}(J_a \cap J_p \neq \phi)$ correctly whp in these cases, the protocol finds $\mathbb{I}(J_a \cap J_p \neq \phi)$ correctly with high probability overall.

We are now close to obtaining a lower bound on the expected number of bits needed for the protocol using the second part of Proposition H.1. Suppose, we gave the worker i access to $|J_a|$. Recall that J_a is a uniformly random subset of J and independent of J_p , conditioned on $|J_a|$. Let the lower bound in Proposition H.1 be $C(\log n)^2$ (i.e., we just named the constant factor C). We prove our result by contradiction. Suppose the protocol requires less than $(C/2)(\log n)^2$ bits in expectation. We found that whp, we have

$$|J_a| \in (0.5n/\log n, 1.5n/\log n).$$

(Notice that these are the same bounds that are needed in Proposition H.1.) Combining this fact and Markov's inequality on the expected number of bits used by the protocol conditioned on $|J_a|$, we deduce that with probability at least $1/3$,⁵⁰ the protocol is faced with problem with $|J_a|$ bounded as required, and such that the expected number of bits the protocol uses for that 'bad' $|J_a|$ is at most $2(C/2)(\log n)^2 = C(\log n)^2$ bits. For each such bad $|J_a|$, the protocol must output something different from $\mathbb{I}(J_a \cap J_p \neq \phi)$ with probability at least ϵ , using Proposition H.1. Since such bad $|J_a|$'s occur with probability at least $1/3$, the overall probability of outputting something different than $\mathbb{I}(J_a \cap J_p \neq \phi)$ is at least $\epsilon/3$. This contradicts that the protocol finds $\mathbb{I}(J_a \cap J_p \neq \phi)$ correctly whp. Thus, we have a contradiction. We conclude that any protocol solving problem P_1 must use at least $(C/2)(\log n)^2$ bits in expectation. Returning to the fact that problem P_1 is at least as hard as finding j who is a stable partner of i with high probability, we conclude that the agent-specific communication complexity for worker i has the same lower bound. Finally, using symmetry over workers (all are in the same tier), we obtain the bound of $\Omega((\log n)^2)$ on the average agent-specific communication complexity. \square

Proof of Theorem 9. This market always has a unique stable matching, which can be constructed by serial dictatorship: workers choose their most preferred unmatched firm in the order of worker tiers. Let i be the worker in question, whose tier is the $(n - \sqrt{n} + 1)$ th from the top. When it is i 's turn to choose under serial dictatorship, there is uniformly random subset of \sqrt{n} unmatched firms remaining. Call this subset of firms J' . Worker i 's unique

⁵⁰We use $1/3$ instead of $1/2$ to accommodate that with small probability, $|J_a|$ may not fall in the desired range.

stable partner is the firm in subset J' that appears highest in her preference list. Let j be i 's unique stable partner. We prove our lower bound by showing that an even easier problem requires $\Omega(\sqrt{n})$ bits of communication: The problem is that of determining, with high probability, if i 's unique stable partner is one of the top \sqrt{n} entries in her preference list. Now this occurs if and only if the set I' consisting of the \sqrt{n} firms that worker i most prefers, intersects with set J' . Note that I' and J' are independent, uniformly random subsets of the set of n firms, each of size \sqrt{n} and I' is known only to agent i whereas J' is known only to the oracle. The lower bound of $\Omega(\sqrt{n})$ follows from first part of Proposition H.1, which is implied by Theorem 8.3 of Babai et al. (1986). Note that for the tier structure described, the same lower bound (up to constant factors) can be obtained for each worker who has rank $n - \Theta(\sqrt{n})$. \square

H.1 Communication complexity of set disjointness

In this section, we prove the a technical result that is used in the proof of Theorem 8.

Consider a set N such that $|N| = n$.⁵¹ Suppose there is a uniformly random subset $A \subset N$ with $|A| = l_a$ known to agent a , and an independent uniformly random subset $B \subset N$ with $|B| = l_b$ known to agent b . Agents a and b are able to interactively communicate with each other and the goal is to determine whether A and B have a nontrivial intersection or not. We are interested in lower bounds for the communication complexity of determining the correct answer with probability of error that vanishes as n grows, although the bounds below hold even for a constant positive error independent of n . Note that the prior probability of intersection between the sets is bounded away from 0 and 1 for any l_a and l_b such that $l_a, l_b = o(n)$ and $l_a \cdot l_b/n \in [0.1, 10]$.

Proposition H.1. *There exists $\epsilon > 0$ such that the following holds:*

- *With $l_a = l_b = \sqrt{n}$, the communication complexity of finding $\mathbb{I}(A \cap B \neq \phi)$ correctly with probability $1 - \epsilon$ is $\Omega(\sqrt{n})$ bits.*
- *With $l_a = \log n$ and $l_b = cn/\log n$ for some $c \in [1/2, 2]$, the communication complexity of finding $\mathbb{I}(A \cap B \neq \phi)$ correctly with probability $1 - \epsilon$ is $\Omega((\log n)^2)$ bits, uniformly over c in the specified range.*

⁵¹We redefine n here. For the purposes of this section there are no workers or firms.

Proof. The first part of the proposition is just the lower-bound part of Theorem 8.3 of Babai et al. (1986).

The second part can be deduced with some work using general results about the communication complexity of disjointness Braverman et al. (2013). We include, a simpler direct proof, inspired by the proof of Babai et al. (1986) for the first part of the proposition. Since the communication setting in the proposition is distributional (i.e. the inputs come from a specified distribution of inputs), it suffices to consider deterministic protocols (since a randomized protocol over a distribution of inputs can be converted into a deterministic one by fixing the random seed that gives the lowest error over the inputs distribution).

Let \mathcal{A} and \mathcal{B} denote the sets of inputs to Alice and Bob. Thus $|\mathcal{A}| = \binom{n}{l_a}$ and $|\mathcal{B}| = \binom{n}{l_b}$. Assume that there is a protocol π of communication cost d . The randomized protocol π induces a partition of $\mathcal{A} \times \mathcal{B}$ into at most 2^d combinatorial rectangles, on each of which the output is either 0 or 1. For all values of c , the probability that the output is 0 (i.e. that the sets are disjoint) is a constant, and therefore, for a sufficiently small ϵ , a constant fraction of the mass is covered by 0-rectangles, on each of which the error rate is $< c_1\epsilon$ for an absolute constant $c_1 > 0$. We will show that the maximum possible mass of each such rectangle is at most $2^{-\Omega(\log^2 n)}$, and therefore there must be at least $2^{\Omega(\log^2 n)}$ such rectangles, and thus $d = \Omega(\log^2 n)$.

Let $R_1 = \mathcal{A}_1 \times \mathcal{B}_1$ be a combinatorial rectangle in $\mathcal{A} \times \mathcal{B}$ such that at most a $(c_1\epsilon)$ -fraction of the elements of R are not disjoint (and thus the 0 output is wrong). We need to show that the size of R is relatively small. Let \mathcal{A}_2 denote the elements in \mathcal{A}_1 that intersect at most a $(2c_1\epsilon)$ -fraction of the elements in \mathcal{B}_1 . Note that we must have $|\mathcal{A}_2| \geq |\mathcal{A}_1|/2$. Let $\mathcal{B}_2 := \mathcal{B}_1$, and $R_2 := \mathcal{A}_2 \times \mathcal{B}_2$. It suffices to show that R_2 has mass at most $2^{-\Omega(\log^2 n)}$.

Construct a sequence of elements S_1, \dots, S_k of \mathcal{A}_2 with the following property: for each i , $|S_i \setminus \cup_{j=1}^{i-1} S_j| \geq l_a/2$. We continue constructing this sequence incrementally until one of two things happens: (1) we cannot add another element to the sequence; or (2) we have $|\cup_{j=1}^k S_j| \geq \sqrt{n}$. We consider each of these cases separately:

Case (1): There are sets S_1, \dots, S_k of \mathcal{A}_2 , such that $|\cup_{j=1}^k S_j| < \sqrt{n}$, and each element $S \in \mathcal{A}_2$ satisfies $|S \setminus \cup_{j=1}^k S_j| \leq l_a/2$. Then this gives the following upper bound on the size of \mathcal{A}_2 :

$$|\mathcal{A}_2| \leq \binom{\sqrt{n}}{l_a/2} \cdot \binom{n}{l_a/2} < n^{-\Omega(l_a)} \cdot \binom{n}{l_a} = 2^{-\Omega(\log^2 n)} \cdot \binom{n}{l_a},$$

and thus \mathcal{A}_2 is small in this case, and we are done.

Case (2): There are sets S_1, \dots, S_k of \mathcal{A}_2 , such that $\sqrt{n} \leq |\cup_{j=1}^k S_j| \leq \sqrt{n} + l_a$, and for each i , $|S_i \setminus \cup_{j=1}^{i-1} S_j| \geq l_a/2$. Each of the S_i 's intersects at most a $(2c_1\epsilon)$ -fraction of the elements in \mathcal{B}_2 , and thus at least half the elements in \mathcal{B}_2 intersect at most $4c_2\epsilon k$ of these sets. Denote these elements in \mathcal{B}_2 by \mathcal{B}_3 . We have $|\mathcal{B}_3| \geq |\mathcal{B}_2|/2$. Each element T in \mathcal{B}_3 can now be described as follows: first specify the S_i 's which T intersects, there are at most $k \cdot \binom{k}{4c_2\epsilon k}$ ways of doing this. Notice that the union of the S_i 's which T does not intersect is at least $(k - 4c_2\epsilon k)l_a/2 > kl_a/3$, since each set contributes at least $l_a/2$ new elements to the union. Therefore, there are at most $\binom{n-kl_a/3}{l_b}$ ways to select the elements of T from the remaining elements. Putting these together we get:

$$\begin{aligned} \frac{|\mathcal{B}_2|}{\binom{n}{l_b}} &\leq \frac{2|\mathcal{B}_3|}{\binom{n}{l_b}} \leq \frac{2k \cdot \binom{k}{4c_2\epsilon k} \cdot \binom{n-kl_a/3}{l_b}}{\binom{n}{l_b}} \leq 2k \cdot \left(\frac{e}{4c_2\epsilon}\right)^{4c_2\epsilon k} \cdot \left(1 - \frac{kl_a}{3n}\right)^{l_b} \\ &\leq_1 2k \cdot e^{k/12} \cdot e^{-kl_al_b/3n} \leq_2 2k \cdot e^{k/12} \cdot e^{-k/6} = 2k \cdot e^{-k/12} \lll_3 2^{-\Omega(\log^2 n)}. \end{aligned}$$

Here \leq_1 holds for a sufficiently small ϵ , since when $4c_2\epsilon < e^{-5}$, $\left(\frac{e}{4c_2\epsilon}\right)^{4c_2\epsilon} < e^{1/12}$; \leq_2 holds because $l_al_b \geq n/2$, and \lll_3 holds because $k > \sqrt{n}/l_a \gg \log^2 n$. Thus \mathcal{B}_2 is very small in this case, and so is R_2 , concluding the proof. \square